Commuting Local Hamiltonians on Expanders, Locally Testable Quantum codes, and the qPCP conjecture

Dorit Aharonov* Lior Eldar†

January 16, 2013

Abstract

Understanding commuting local Hamiltonians (CLHs) is at the heart of many questions in quantum computational complexity and quantum physics: quantum error correcting codes, quantum NP, the PCP conjecture, topological order and more.

We first consider the complexity of the problem of approximating the ground value of a k-local CLH; The complexity of the exact case remains unclear despite significant effort [18, 3, 51, 37], and the approximation case is tightly related to the major conjecture of quantum PCP. We show that if the underlying interacting graph of the CLH instance has small-set expansion which is ε close to optimal, then the approximation problem to within a factor $O(\varepsilon)$ lies in NP. Thus, the better the expander, the easier it is to approximate. This puts a bound on the complexity of CLH on small-set expanders, and indicates that known PCP constructions such as Dinur's [27], whose output instances are small-set expanders, cannot be used to prove a PCP theorem for CLH.

We next proceed to study the much related topic of locally testable quantum codes based on local stabilizers; these are ground states of CLHs. We show that the better the small-set expansion of the interaction graph underlying the stabilizer becomes, the *less* robust the code becomes. This phenomenon seems to be inherently quantum. We combine the above result with an upper bound we prove on the robustness of stabilizer codes with *bad* expansion (which is a classical phenomenon) to derive a bound on the robustness of quantum LTCs with an *arbitrary* underlying interaction graphs. We derive non-trivial bounds on the robustness of quantum LTC codes; To the best of our knowledge the quantum LTCs were not studied before.

Finally, we initiate the study of the quantum analogue of PCPs of proximity [13] which play a central role in classical PCPs; it is known that PCPPs induce LTCs with related parameters [13], and we show this holds also in the quantum world; this connection between LTCs and PCPs further motivates the study of LTCs and their possible parameters. We defer the exposition of this part of our work to the next version of the paper.

Much is left for further research, and in particular improving the range or applicability and generality of our results. Nevertheless we believe that the results highlight interesting limitations on quantum PCP and that they point at possible routes that might shed light on the stubborn PCP conjecture.

*School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel
†School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel.

1 Introduction

This paper is concerned with commuting local Hamiltonians (CLHs). A k-local Hamiltonian is a Hermitian matrix $H = \sum_i H_i$ operating on the Hilbert space of n d-dimensional particles, where each term H_i (sometimes called constraint) acts non-trivially on at most k particles. For the Hamiltonian to be a k-local *commuting* Hamiltonian we require in addition that every two terms H_i, H_j commute; W.L.O.G we assume in this case that H_i's are projections. We denote the family of such Hamiltonians by $CLH(k, d)$. Commutative local Hamiltonians have been the subject of very intense research for many years in physics; As it turns out, CLHs are related to some of the deepest questions in quantum computational complexity as well. Let us first provide the background and context for the questions we ask here regarding CLHs.

1.1 Background and Context

The computational complexity of the Local Hamiltonian problem Kitaev first defined in 1998 the local Hamiltonian (LH) problem, where one is given a local Hamiltonian on n qudits and two real numbers a, b, with $b-a \geq 1/poly(n)$, and is asked whether the lowest eigenvalue is at most a or larger than b; He showed [42] that this problem is complete for the class QMA, the quantum analogue of NP, providing a quantum counterpart of the celebrated Cook-Levin theorem.

In 2003, Bravyi and Vyalyi [18] considered the similar decision problem for the commuting case; The commuting restriction might seem at first sight to devoid the LH problem of its quantum nature, since all terms can be diagonalized together, and moreover, one usually attributes the interesting features of quantum mechanics to its non-commutative nature, cf. the Heizenberg uncertainty principle; following this one might suspect that the commuting version trivially belongs to NP. This intuition however is misleading, since ground states of CLHs can exhibit intricate entanglement phenomenon, such as the topological order exhibited by the Toric code [43]. Using clever applications of representations of C^*-algebras Bravyi and Vyalyi [18] gave a proof that the two-local commuting case indeed lies within NP, and this was generalized later to more general Hamiltonian classes [3, 37, 51]; However the complexity of the general CLH problem remains unresolved.

The resolution of the complexity of the CLH problem seems important on its own right, given the importance of commuting Hamiltonians in physics, its intriguing nature, its relation to multiparticle entanglement and its connection to topological order. But apart from these reasons, it is also related to the major open problem of whether a quantum analogue of the celebrated PCP (Probabilistically Checkable Proof) theorem holds; we will discuss this slightly later.

CLH and quantum Error correcting codes Arguably the most important class of quantum error correcting codes are the stabilizer codes (see [34]); A stabilizer code is defined by a set of commuting generators, each generator being an element of the (generalized) Pauli group on n d-dimensional particles, i.e. a tensor product of n Pauli operators. The code is defined to be the simultaneous 1 eigenspace of all generators (therefore these generators are called *stabilizers*); if all stabilizers are k-local (namely, only k out of the n Pauli's are non-identity), this can easily be seen to be equivalent to a ground space of a $CLH(k, d)$ instance. Thus, CLHs and quantum stabilizer codes are intimately connected. A particularly beautiful example is that of the Toric code of Kitaev [43] mentioned before, due to the non-local entanglement it exhibits; Various other stabilizer codes exhibiting topological order were defined (see e.g. quantum double models [45] and many other stabilizer codes were con-

structed and studied (eg, [34, 53, 54, 23, 25, 50, 58]). Many questions regarding stabilizer codes are of interest with much left to be understood: what are the possible parameters (distance, rate, etc.) of such codes, and what are the tradeoffs between those parameters [21, 31, 59], how efficient can the description of the states in those codes be [2], what are the various properties of such codes with respect to perturbations [19, 20] and thermal fluctuations [56], their behavior in the presence of random noise [26, 46], all the above when constraints are restricted to be local, namely, the quantum analogue of LDPC codes ([58, 59, 46, 47] and references therein), and more. All those questions can be viewed as questions on CLHs, and their resolution is tightly related to our understanding of multiparticle entanglement, and of course to our understanding of the capabilities of quantum error correction.

Quantum PCP Both the above topics are related to what is now considered a major challenge in quantum complexity theory: the quantum PCP conjecture. The classical PCP theorem [8, 7, 27] is arguably the most important discovery in classical theoretical computer science over the past two decades; it states that there is a polynomial time reduction that maps a given Constraint Satisfaction Problem (CSP) instance (i.e., a collection of k-local constraints on n Boolean or d-state variables) into another CSP instance, such that if the original instance is satisfiable, then so is its image, but if the original instance is not satisfiable, then no assignment satisfies more than some constant $c < 1$ fraction of constraints in the image instance. This implies that it is possible to check whether an assignment satisfies a given CSP or not, by reading only a constant number of random locations in the assignment (!). A major open problem is whether some quantum version of the PCP theorem holds [4, 1]. A quantum PCP conjecture can be phrased as follows:

Conjecture 1 Quantum PCP (qPCP) *There exist constants $c > 0, k$, and a (quantum) poly-time algorithm, which takes an instance of a k-local Hamiltonian H, to another instance of a k-local Hamiltonian H', such that if H is satisfiable (namely, has ground energy 0), then so is H', but if H is unsatisfiable, the minimal eigenvalue of H' is at least $c||H'||$.*

A different way to state the theorem is to say that even the approximation of the ground energy of the given Hamiltonian to within a constant times its norm, is quantum NP hard; an equivalent formulation (see [4] for exact statement and a proof) is that QMA-hard problems have a formulation in which the quantum witness can be tested by measuring only a constant number of random qudits.

Attempts to extend the classical proofs of the PCP theorem to the quantum settings, or to disprove it, have encountered severe obstacles so far [4, 6, 37, 38]. In particular, the classical proofs rely heavily on copying classical information; we cannot do the same in the quantum world, due to the no-cloning theorem. It seems that we do not understand some crucial issue about the local versus global behaviors of entanglement which is central to the problem. The resolution of the quantum PCP conjecture, either to the positive or the negative, is likely to shed light on multiparticle entanglement, with implications to computational complexity as well as to physics, e.g., to hardness of approximation in the quantum setting, robustness of quantum correlations at room temperature, quantum error correction, topological order and more.

Given the difficulty in making progress on the quantum PCP question so far, it is natural to restrict attention to commuting Hamiltonians. One one hand, one might hope to prove a quantum PCP theorem for commuting Hamiltonians, with the hope of borrowing ideas from such a proof to the more general case (this route might be useful even if eventually the general CLH problem turns out to be inside NP, which would collapse the commuting PCP theorem into the already known classical

PCP theorem). On the other hand, upper bounds on the hardness of approximating the CLH problem (such as putting the CLH approximation problem inside NP for certain approximation factors) could be considered as initial steps towards disproving the general PCP conjecture, or at least as limitations to its strength. This approach of attacking first the commuting case was taken by Aharonov and Eldar [3], Hastings [37], and Hastings and Freedman [38].

Locally Testable codes Classical PCP proofs rely heavily on a particular type of error correcting codes, called LTC codes (Locally Testable Codes) which are interesting in their own right. These are codes which satisfy the following property: Given a word of fractional distance $\delta > 0$ from the code, we would like to detect that it is not in the code by randomly choosing a constant number k out of the n locations, and performing a simple test on those locations. The code is called a locally testable code (LTC) if such a test succeeds with probability $s > 0$, given that $\delta > 0$ (there are variants on this definition, which we will discuss below). Note that the requirement for soundness when querying a constant number of locations is reminiscent of the similar goal in the context of PCP, in which querying a constant number of bits of a proof is supposed to detect that the proof is incorrect with constant probability. Indeed, locally testable codes, such as the Hadamard and the long code [36, 16] have played a central role in the proofs of the PCP theorem ([8, 7, 27]).

One can make the connection between the classical PCP theorem and LTCs more rigorous. It is possible to define a strong version of PCP, called PCP of proximity (PCPP) [13], which is a structure hidden in most known PCP proofs; Ben-Sasson et. al. showed [13] how given a PCPP, one can take a good code and construct from it a good LTC code, which inherits its local testability parameters (soundness and number of queries) from those of the PCPP [13].

To the best of our knowledge, quantum LTCs and their robustness were not studied before, and neither was the notion of quantum PCPPs.

1.2 Our contribution

In this paper we attempt to gain new insights into the above topics. To do that, we consider an important facet, central to all the above issues: the topology of the interaction graph underlying the set of constraints.

Constraint satisfaction problems on a lattice, both in the quantum and classical case, can be easily approximated in polynomial time to within factors which are even sub-constant. This is done by throwing away constraints that disconnect the lattice to logarithmic sized boxes which can be solved separately. On the other hand when the underlying graphs of the classical problem are expanders, the constraint satisfaction problem can exhibit extreme robustness against even constant approximation factors, as in the PCP theorem (see Dinur's proof [27, 9]).

Expansion of the interaction graph plays a crucial role in the context of error correcting codes as well; Classical expander codes were defined in [57, 5, 52]. In the quantum regime quantum codes on expanders have not been explored very thoroughly, though see [58, 38].

In the above mentioned results, expansion is referred to loosely; several definitions of expansion are used in the literature, and we have not specified which one is being used. To be concrete, we need to choose a good expansion definition for our context. We first observe that crucially, the graphs we work with are k-local hypergraphs for $k \geq 3$. This because the two-local case of commuting Hamiltonians is not interesting from our point of view; 2-local CLH is known to be in NP [18], and exhibits

only local entanglement, and stabilizer codes with 2-local check terms cannot correct quantum errors for similar reasons. We therefore need to work with hypergraphs, or with CLHs with $k \geq 3$ locality.

One possible route to take is to use one of the known definitions for expansion in hypergraphs: geometric, homological or spectral; (for references and a survey see [48]). However it is unclear which of those makes more sense and the relations between them are not all known [48]. One could also choose to work with a graph underlying the hypergraph (connecting any two nodes that appear in the same constraint); This definition loses much of the "structure" of the original hypergraph. We choose to work here with the natural bi-partite graph induced by the hypergraph (as in e.g., [52, 24]). We define the induced bi-partite graph $G = (L, R; E)$ with constraints on the left, and variables / qudits on the right (see definition 2), and a variable is connected to the constraints it appears in. We say that such a graph is an ε small-set bi-partite expander, if for any set S of size at most k particles, the number of local terms incident on these particles is at least $D_R|S|(1-\varepsilon)$, where D_R is the right degree of the graph. We note that $D_R|S|$ is the maximal possible number of constraints acting on those particles, so ε can be viewed a "correction" to this number.

We remark that our definition cares about the expansion of only constant-size sets, whereas usually, one refers to *small-set expanders* as those graphs where all sets of size some linear fraction of $|R|$ are required to expand [29, 49]. We will address this point also in section (1.5), following the exposition of our results.

We make use of a very simple property of very good small-set expanders, namely that most of the neighbors of a set of size k only touch this set at one point (see Fact 1 in Subsection 1.3). This turns out to be an extremely useful property which will bare strong consequences for both the computational complexity of CLHs on small-set expanders, as well as for stabilizers whose graphs are small-set expanders.

1.2.1 Approximating CLHs on expanders

We start by studying the complexity of the approximate version of the CLH problem; we ask whether we can bound from above the complexity of CLH if we are allowed a constant error in the minimal energy; Specifically, if we allow to throw away a certain constant fraction of the constraints; For that matter, we study what seem to be "hardest" class of interaction graphs, namely small set expanders.

We show that the approximation version of the $CLH(k,d)$ problem becomes "easy" and falls into NP. (This is of course more surprising if one believes that the general CLH problem is not in NP; if on the other hand one believes CLH is in NP, this can be considered as a step towards this goal.)

Our first theorem states that:

Theorem 1 *(the approximation of CLH on small set bi-partite expanders is in NP) Let $\gamma(\varepsilon) = 2kd\varepsilon$. Let H be an instance of $CLH(k,d)$ for constant k, d whose bi-partite interaction graph has a right degree D_R, and is ε-small-set expanding, for $\varepsilon < \frac{1}{2}$. Let $\lambda_1 > \lambda_2 > \ldots > \lambda_N \geq 0$ be the eigenvalues of H. Then for every eigenspace λ_i, there exists a state $|\psi_i\rangle$ such that $|\|H|\psi_i\rangle\| - \lambda_i| \leq \gamma(\varepsilon)\|H\|$, and such that $|\psi_i\rangle$ can be generated by a constant depth quantum circuit. In particular, the $\gamma(\varepsilon)$-approximation problem of $CLH(k,d)$ on such ε small-set bi-partite expander graphs is in NP.*

For the above theorem to be non-trivial, ε must be sufficiently small, so that $\gamma(\varepsilon) = 2kd\varepsilon$ is less than 1.

One might ask whether Theorem 1 is trivially true; we might have allowed to remove enough terms so that the graph is disconnected into small components and the problem becomes solvable in P, just like what happens in the lattice case. We show that this is not the case:

Theorem 2 *Given $\varepsilon < 1/2$, it is NP-hard to approximate $CLH(k, d)$ whose bi-partite interaction graph is ε small-set expanding, to within a factor $\gamma(\varepsilon) = 2kd\varepsilon$.*

Theorem 1 can be considered as a first step towards proving that the CLH problem lies inside NP;

Alternatively, if one takes the point of view that CLH is not inside NP, and hence that a quantum PCP theorem might in fact be proven via showing hardness of approximations of CLHs (as is the approach of [37, 38]), then Theorem 1 puts limitations on such a qPCP. While standard constructions of classical PCP use graphs with excellent small-set expansion (see the proof of Theorem 2) and one can make these graphs have arbitrarily good small-set expansion (albeit while increasing the degree), Theorem 1 implies that the smaller the expansion error of the graphs underlying the CLHs are, the weaker the qPCP theorem using these CLHs becomes. The approximation error for which hardness is shown is bounded from *above* by some constant times the expansion error. This contradicts classical intution. We discuss this further in the Discussion, Subsection 1.5.

1.2.2 Quantum stabilizer Locally-Testable Codes

Next, we study the much related topic of stabilizer codes. In particular, we study how stabilizer codes with local stabilizers behave in the context of local testability.

As mentioned above, classical locally testable codes satisfy the following property: Given a word which is of relative distance δ from the code, then if δ is not 0 then this fact can be detected with some constant probability $s > 0$ by querying a *constant* number k (hence, *locally* testable) of coordinates in the code; One can consider strong LTC codes in which the probability s is proportional to δ, even for sub constant δ's, or weak LTC codes in which the probability s is at least a constant given that δ is at least a constant. The essential parameter of LTC codes is thus its robustness, captured roughly by how s relates to δ.

More precisely, we say that a locally testable code with degree D_R (namely, that each qudit is examined by D_R constraints) is $r(\delta)$-robust if an error on δn locations violates a at least $r(\delta)D_R\delta n$ constraints, namely, a fraction $r(\delta)$ of the maximal number it can possibly violate (See Definitions 8, 9 for a rigorous definition of the weight of an error and the robustness for various weights).

There are two famous classical LTC codes used in the PCP proofs of [7, 27, 9]: the long code and Hadamard codes. They exhibit extreme robustness, in the following sense: for any $\varepsilon > 0$, there exists $\delta > 0$ such that all error patterns of weight at most δn (for code-length n) are violated by a fraction at least $1 - \varepsilon$ of their incident (3-local) constraints).

We would like to study the robustness of quantum error correcting codes; To the best of our knowledge, the question of local testability of Quantum error correcting codes was not studied yet. The very notion of weight of an error is not very natural in the context of general quantum error correcting codes; we leave the general definition to a later version of this paper. For stabilizer codes, however, the notion of weight of an error is very natural: Just count the number of non-identity Pauli's in the error (essentially, modulo the stabilizer group). For the rigorous definition see 8. Robustness would thus correspond to how many stabilizers will not commute with such an error.

The robustness of quantum stabilizer codes with local stabilizers seems, on the face of it, to be inherently restricted. It is illuminating to consider in this context the Toric code example [43]. The Toric

Code is defined by a set of 4-local constraints on a lattice of $\sqrt{n}$ by $\sqrt{n}$ qubits. It is well known that one can consider a chain of errors of length $\theta(\sqrt{n})$ of qubits, such that the only violated constraints will be those two constraints that touch the end qubits of this chain. This already indicates very low robustness ($r(\delta) < \frac{1}{\sqrt{n}}$) for values of δ of the order of $\frac{1}{\sqrt{n}}$. It is easy to extend the low robustness behavior to words of linear distance from the code, by scattering short chains of errors of length $\theta(1/\varepsilon)$, which are separated from each other by at least $\theta(1/\varepsilon)$ lattice sites.

The message of this part of our paper is that quantum stabilizer codes in general, as long as they have local stabilizers, exhibit inherent non-trivial (and inherently quantum) upper bounds on their robustness.

Our first result in this context might be viewed as surprising, as once again it poses constraints when the underlying interaction is expanding: it turns out that if the underlying graph is a good small set bi-partite expander, then the robustness is as low as the expansion error itself, for all errors with weight smaller than some constant δ.

Theorem 3 *Let C be a stabilizer code, with minimal distance > 1, and a k-local generating set $\mathcal{G} \subset \Pi^n$, such that each qudit is examined by D_R generators. Suppose the bi-partite interaction graph of $\mathcal{G}$ is ε-small set bi-partite expander, for $\varepsilon < 1/2$. Then, for all $\delta < \min\{\frac{1}{k^3 D_R}, \frac{1}{2} dist(C)\}$, we have $r(\delta) \leq 2\varepsilon$.*

In the above $dist(C)$ denotes the distance of the code (see Definition 4). Theorem 3 thus provides a trade-off between the expansion and the robustness; The better the expansion, the worse the robustness. This seems like an inherently quantum phenomenon, contrary to what happens in the classical world.

To contrast it with the classical setting, recall a result by Dinur and Kaufmann which shows that a robust LTC code must have an interaction graph which is a small-set expander (Theorem 1.1 in [28]); We note that direct comparison does not hold since the definitions of expansions used are different; [28] does not use bi-partite graph expansion but rather the graph in which an edge connects any two nodes that participate in a common constraint.

Stronger evidence that Theorem 3 shows a solely quantum phenomenon is given by the construction of *lossless expanders* in [24]. This construction gives classical error correcting codes whose robustness is very close to 1.

Claim 1 *For any $\varepsilon \in (0, 1/2)$, and $r \in (0, 1)$ there exists a constant $\delta = \delta(r, \varepsilon)$, such that there exists an explicit infinite family of codes $\{C_\varepsilon(n)\}_{n \in \mathbf{N}}$, of n bits, of constant fractional rate r, and constant fractional distance $d = d(\varepsilon, r)$, whose check terms are of locality whose expectation is equal to a constant k, and all errors of weight less than δn have robustness at least $1 - 3\varepsilon$.*

The proof of this claim is given in the appendix. We note that contrasting Claim 1 with Theorem 3 is not exact, since the classical codes constructed have only average constant degree and not constant degree; we believe this is not an essential point.

We use theorem 3 together with a rather involved probabilistic argument, to establish an *absolute* non-trivial upper-bound on robustness of stabilizer codes, for errors of size at most some constant relative distance. A semi-trivial such bound exists, due to the size of the alphabet: For a given qudit of dimension d, the number of possible errors on one qudit is $d^2 - 1$. Fix a qudit, and pick Q to be the Pauli on that qudit which is most popular among all the stabilizers acting on that qudit. Then at least $1/(d^2 - 1)$ of the stabilizers on that qudit will commute with that error, and so at most $\alpha(d) = 1 - \frac{1}{d^2 - 1}$

of the stabilizers touching that qudit will not commute with that error, and thus will detect the error (see fact 4 for an exact proof). The robustness is thus bounded by $\alpha(d)$, which we call the *single-error robustness*, or the alphabetical upper bound on the robustness. Classically, there is no direct analogue to the requirement of non-commutativity to achieve constraint violation, and hence $\alpha(d) = 1$ in the classical case.

We are interested in a deeper phenomenon; it turns out that for any k-local stabilizer codes, the robustness must be strictly smaller than the *single-error robustness*.

Theorem 4 *(Roughly) For any stabilizer code C of k-local terms ($k \geq 4$) over d-dimensional qudits, where each qudit interacts with D_R local terms, errors of weight at most some constant fraction of n have robustness at most $\alpha(d)(1 - \gamma_{gap})$ for some $\gamma_{gap} = \gamma_{gap}(k, d) > 0$.*

This indicates some inherent "problem" in the quantum setting when one is interested in maximizing the robustness to make the code most 'sensitive" (in terms of energy penalty) to errors; Understanding this phenomenon might be tightly related to the clarification of the qPCP conjecture.

1.2.3 PCPs of Proximity

A PCP of proximity, or $PCPP$, is a form of language verification where the verifier receives, in addition to the usual NP-witness to the veracity of the claim made by the prover, an auxiliary proof that allows for a low-query test on whether or not the witness is "close" to or "far" from the set of acceptable witnesses. Ben Sasson et al [13] provided construction of an LTC code given a PCPP, and described how the parameters are mapped from one to another. (see Construction 4.3, and Proposition 4.4 in [13]).

We define quantum PCPPs, where the PCPPs must be defined with respect to a given set of constraints that can be used. We then show that a similar result to that of [13] mapping a quantum PCPP to an LTC with a corresponding set of constraints, holds also in the quantum setting. Thus, limitations on quantum LTC codes translate to limitations on quantum PCPPs using similar sets of constraints; providing an important motivation to understand Quantum LTC codes.

We will delay the definition of quantum PCPP using various sets of constraints, together with the general definition of LTCs (which use not only stabilizer constraints) and the above statement connecting the two, to the next version of this paper.

We mention that for now, our bounds on robustness of LTC codes are merely initial indications of the difference between the quantum and the classical LTC behavior, and much is left for improvement in terms of the parameters; moreover, currently our bounds on the robustness of LTCs hold only for stabilizer codes. Thus, they imply rather limited implications on qPCPPs. However, the connection between qLTCs and qPCPPs extending the known classical connection between LTCs and PCPPs strongly motivates further study of qLTCs, as well as improving our parameters and extending our results to other constraints systems; a deeper study of the implications to the qPCP conjecture of this direction is called for.

1.3 Overview of Proofs

1.3.1 Simple observations on expanders

Given a bi-partite graph $G(R, L : E)$, we say a set of nodes in R is ε-expanding if $\Gamma(S) \geq D_R|S|(1-\varepsilon)$ where $\Gamma(S)$ the neighbors of S and D_R the right degree.

Two basic observations in this paper are the following; The proofs can be found in the appendix. (see [24] for similar observations)

Fact 1 *Consider $S \subseteq R$ in a bi-partite graph $G(R, L : E)$ and let S ε-expanding, for $\varepsilon < \frac{1}{2}$. Then a fraction at most 2ε of all vertices of $\Gamma(S)$ have degree strictly larger than 1 in S.*

Fact 2 *Let $S \subseteq R$ in a bi-partite graph $G = (R, L; E)$, such that S is ε expanding. Then there exists a vertex $q \in S$, such that the fraction of neighbors of q with at least two neighbors in S is at most 2ε.*

1.3.2 The complexity of approximating CLH on expanders

The idea behind the proof of Theorem (1) is as follows. Much of our current knowledge on the complexity of CLH relies on a critical observation by [18] that any two commuting local terms H_i, H_j can be viewed as acting on "disjoint" subsystems, using some basic facts from the representation theory of $C*$-algebra. Informally (see Subsection 2.2 for a formal description), the idea is this. We consider two terms in the Hamiltonian, H_i and H_j, which intersect on some subset of qudits, whose Hilbert space is $\mathcal{H}_{int}$. Bravyi and Vyalyi's lemma claims that there exists a local isometry on $\mathcal{H}_{int}$, preserved by H_i, H_j, which if applied, the space of the intersection can be written as a direct sum of subspaces, and in each one, the two terms are "disconnected" from each other, so that they essentially act on different subsystems which are in tensor product. In other words,

$$\mathcal{H}_{int} = \bigoplus_{\alpha} \mathcal{H}_{int}^{\alpha,i} \otimes \mathcal{H}_{int}^{\alpha,j} \tag{1}$$

Such that H_i (H_j) restricted to the α subspace acts non trivially only on the left (right) subsystem, $\mathcal{H}_{int}^{\alpha,i}$ ($\mathcal{H}_{int}^{\alpha,j}$). This lemma was used by [18] to show that 2-local CLHs (corresponding to graphs) are in NP; However as was noted in [3] more is needed in the case of hypergraphs.

In the context of expanders, we use it as follows. Given a constraint H_i, if we now remove all terms that share at least 2 vertices with H_i, (and this, by Claim 1 is just a small portion of the terms intersecting that term) we get into a situation in which all terms that intersect H_i, intersect it at only one particle. This makes the intersection pattern simple enough, so that Bravyi and Vyalyi's lemma [18] can be applied almost directly on each of the qudits on which the terms acts on. The final result is that the term H_i can be separated from the remaining terms and can be viewed (up to the relevant isometries) as acting on a separated subsystem.

How to use this insight to show that the approximation to within a small constant is in NP? It is not enough to simply take away, for each term, all terms that make it isolated, as this will require removing too many terms. Instead, we use the above idea iteratively. At each iteration we choose a local term v and "isolate" it (see definition 7): we remove as many terms as necessary so that among the remaining terms, no local term shares more than one particle with v (see figure 2). Bravyi and Vyalyi's lemma can then be applied, and v can be separated from the rest, after restricting the qudits to the relevant subspaces in their direct sum of Equation 1. Now we can continue in the same way, starting from the restricted subspace. This way we gradually "tear-away" local subspaces of particles.

The analysis of the upper bound on the number of terms that need to be removed altogether requires some thought, since the number of particles does not necessarily decrease through one iteration; it is only their dimensionality that decreases. We resort to an amortized analysis counting the number of *dimensions* that are removed in total. As far as we know this amortized counting of

dimensions is novel in the context of quantum Hamiltonian complexity, and may be of interest on its own.

Note that unlike in the lattice case, removing edges here is done not in order to *disconnect* part of a graph, but only to isolate terms; loosely speaking, to make the interaction pattern more sparse. Disconnecting entire parts of the graph will require removing too many terms in the context of expanders, and will thus lead to too large an approximation error. The proof of theorem (2) that the approximation problem on small-set expanders is NP-hard follows from the observation that the output of the PCP amplification routine due to Dinur ([27]) is an excellent small-set expander.

1.3.3 Bounds on LTC codes on Expanders

To prove Theorem 3, we argue that a stabilizer code whose underlying interaction graph is a good enough bi-partite small-set expander, cannot be too robust. To do this, we define a particular error pattern, which has a large weight modulo the stabilizer, but which does not violate too many generators.

The error pattern we define acts on one qubit in each one of a set of "isolated" terms. By "isolated" term (formally, L-independent terms - see Definition 11), we mean something completely different than the previously used isolated qudits; a set of terms is L-indepdent if no two terms in it are closer than some constant in the interaction graph. Picking one qudit in each term, we claim that the weight of this error cannot decrease when multiplied by an element from the group. Intuitively, this is because of the following reason; If such an error pattern can be represented more succinctly, there must be one single qudit whose error is removed modulo the stabilizer group, and moveover, no new errors appear (modulo the group, again) in a constant but large enough neighborhood of that qudit. We show that this cannot happen since the stabilizers are too local.

We need to design an error pattern whose robustness will be limited. By examining the code itself, we can "distill" an error which causes a minimal number of violated constraints. Here Fact 2 allows to pick cleverly the qudit in each term on which the error will act, and also to define what Pauli error will act on it to minimize the penalty (namely, the number of generators which will not commute with the error). For a term g we choose the qudit q to be the one promised by the lemma: a qudit s.t. a vast majority of the terms acting on it intersects g at only one location. We define the error on q to be the restriction of g to q; since all stabilizer terms commute with g, those that intersect g only on q must agree with it on q. It follows that most of the incident generators on q will commute with this error, and only a meager fraction of its incident generators, will realize that this is indeed an error.

1.3.4 Upper bound on robustness

Our final theorem states that for any constants $k \geq 4, d$, there exists a constant less than 1, upper-bounding the robustness of any quantum LTC of k-local constraints, on d-dimension particles. Our proof relies on bounding the robustness of a quantum LTC from two sides. On one hand, we use the bound of theorem (3) which is, in fact, the "surprising" side, which implies that high expansion forces low robustness. We then add a new claim showing that quantum Stabilizers, not only suffer from the quantum effect of (3) but also, cannot avoid, the "classical" effect that codes with poor expansion have low robustness.

We restrict our attention to sets of qudits belonging to generators which are "far" away from each other in the interaction graph of the generating set. Given a set of qudit "islands" with poor collective

expansion, we try to find an error on this set which has less-than-optimal number of violations. At least intuitively, a set with poor expansion, cannot have full robustness, since if "all" bits of the set are erroneous there are less than an optimal number of constraints examining this set, so there would be less-than-optimal number of violations. However, this intuition on its own is insufficient, since we don't know whether or not we can find an error whose weight will be the entire set. Rather, we resort to using a random error, so that only a small fraction of the set is erred. We show, as a sub-claim that on average, even a random error which is "sparse" on the set, will still sense the fact that the complete set has sub-optimal expansion, and would yield less than optimal number of violations.

The fact that the resulting error has large weight modulo the centralizer follows from an observation we prove, that each "island" of qudits cannot have its errors "erased" via a more succinct representation, if most of the "island" is uninhabited, i.e. less than $1/2$ of the qudits are erred. We then have a fine trade-off: on one hand, we would like the number of errors on a given "island" to be smaller than $1/2$, and on the other hand, we would like a good portion of "islands" to have at least 2 errors, so that they "sense" the sub-optimal expansion.

1.4 Comparison with prior work

The question of approximating the minimal energy of local Hamiltonians was considered by various researchers; see Bansai, Bravyi, Terhal [11], and Kempe and Gharibian ([33]), who gave upper bounds on non-trivial approximation factors for general local Hamiltonians.

The approximation of commuting local Hamiltonians, and the PCP conjecture in general, was the motivation of Hasting's work [37] in which he showed that certain classes of CLHs lie in NP (without approximation). Arad [6] also considered the commuting case as a base for perturbations in his partial no-go for quantum PCP result. Hastings and Freedman [38] recently defined generalization of expanders to hypergraphs, motivated by the attempt to prove the quantum PCP conjecture. We note that the graphs designed in [38] are small-set expanders to which our theorems, and in particular Theorem 1 apply; hence, our work puts limitations on the strength of quantum PCPs using such graphs.

In a recent independent result, Brandão and Harrow [17] have shown that the approximation of 2-local Hamiltonians on expanding graphs (here expanding in the standard definition) lies in NP, with an error depending on the second eigenvalue of the interaction graph. This is a complementary result to our theorem 1, since on one hand, there is no commutation requirement in [17], but on the other hand they only handle the case of graphs, while higher locality might be crucial in the context of the problems studied here. In particular, it is unclear whether it is possible to derive a result about approximating k-local Hamiltonians on expanders from the results of [17], since the gadgets allowing to move from $k-$ to 2-local [22] change the topology of the graph.

We note that our Theorem (1) reduces in 2-local case to a 0 error approximation, namely an exact verification procedure; in other words, it reduces to the result of ([18]). This is because bi-partite expansion is essentially maximal for any 2-local CLH, since in this case two different constraints cannot intersect on more than on qudit. This is another indication that at least for the CLH case, bi-Partite expansion (Definition (3)) seems to be the appropriate one to use, compared to that of regular expansion, which is used by [17] for the general LH case.

As for quantum error-correcting codes, quantum LDPC (Low Density Parity Check) codes, namely codes with local check terms, were already studied in the literature (see [58, 59] and references therein, as well as [46, 47]). However, as mentioned before, we are not aware of any prior work

on quantum locally-testable codes.

1.5 Discussion

Complexity of CLH One might hope that Theorem 1 could be extended to show that the approximation of the CLH problem on a *general* topology is in NP, perhaps by bridging between our results on expanders to the easy case of approximation on lattices. We note that a hint about the difficulty of such a result comes from the status of the Unique-Game-Conjecture [41] which is known to be easy on both cases (albeit for interaction graphs and not hypergraphs), though conjectured to be NP-hard in general. Still, it may be possible to achieve some weaker approximation of general graph using "bridging" between the case of lattices and expanders, perhaps by sub-exponential witnesses, following the work of [55],[10]. We note here that the "bridging" that we did achieve for stabilizer code in Theorem (4), indeed depends on these instances being more structured, and as such, does not follow through for general CLH instances. It thus remains an important open problem to generalize our results to all interaction graphs. Of course, clarifying the exact case remains a major open problem.

Quantum PCP and quantum LTCs Though one cannot derive using Theorem (1) a bound for general PCP, it does place severe restrictions on any commutative version of a quantum PCP: i.e. a poly-time algorithm taking any instance of LH (or even just CLH), to a CLH instance with constant promise gap. Specifically, assuming CLH is not contained in NP, our results indicate that the output of such a quantum PCP is unlikely to assume the form of a bi-partite small-set expander, unlike in the classical PCP proofs. First, Theorem 1 implies that for a given expansion error ε, the promise gap must be of the order of ε which can be a very small constant. Moreover, one may speculate that if the output of such a PCP routine, has some small expansion error, than this error may be set arbitrarily low. This would imply that as ε goes to 0, the promise gap goes to 0, contradicting the existence of such PCPs.

We note that to the best of our knowledge, the question of general upper bounds on robustness of quantum LTCs was not considered before, though of course robustness of specific quantum codes is very well understood; interestingly, the prominent example of the Toric codes has inherently small robustness, which is tightly related to the fact that the code exhibits topological order, and thus only the "end points" of an error matter [42]. Quantum error correcting codes can be viewed as a generalized topological order (see [37]), and perhaps some bounds on robustness phenomena can be carried over to all codes, not necessarily stabilizer LTCs as in our Theorem 4, perhaps by using the fact that the code has large distance; It will be very interesting to derive such a result. Also, it will be most interesting to come up with explicit (or even non-explicit) constructions of quantum codes with very good robustness.

Of course, strengthening of our results in various other directions are called for: be it generalizing to all weights of errors, to various types of constraints, and of course, removing the commuting requirement.

In a later version we will define Quantum PCPP; and make explicit the connection between quantum PCPP and quantum LTCs. In particular, this implies that the restrictions on LTC codes with various types of constraints lead to bounds on PCPPs using similar constraints. The results we have presented here (3) and (4) provide unexpected limitation on stabilizer LTC codes. The connection to quantum PCPP strongly motivates extensions and further clarifications of our results on qLTCs and their implications to qPCPs.

Finally, we believe the results we presented here hint that possibly, the problem of CLH and clarification of the importance of the commutation relations, might be more relevant to the resolution of the quantum PCP conjecture than was thought before. For example, if the qPCP conjecture were true, and the method of proof was indeed through a Quantum PCPP, and if indeed CLH is not Quantum NP hard (all of the above conjectures might seem quite reasonable to believe in at this point), then this must imply the existence of novel excellent qLTC codes whose local check terms are non-commuting, by which such a qPCP theorem would be proven. Can such codes exist? Studying this question directly can shed light on the likelihood of the above mentioned conjectures to be true.

Organization of paper

In Section 2 we provide the background. Section 3.1 proves theorem 1, showing that approximation of CLH on expanders lies in NP. Section 3.2 proves Theorem 2, showing that the approximation problem we handle is NP-hard. Section 4 provides bounds on the robustness of quantum LTC codes on expanders, and Section 5 provides an absolute bound on robustness of stabilizer LTC codes regardless of the expansion of their underlying graph.

2 Background

A (k,d)-local Hamiltonian on n qudits is a Hermitian matrix $H = \sum_i H_i$ operating on a Hilbert space $\mathcal{H} = \mathbf{C}^{d^{\otimes n}}$ of n d-dimensional particles, where we assume for all i, $\|H_i\| = 1$ and $H_i = h_i \otimes I_{n-k}$ is Hermitian operating non-trivially only on k particles.

Definition 1 ***The commuting local Hamiltonian problem*** *In the (k,d)-local Hamiltonian problem on n d-dimensional qudits we are given a (k,d)-local Hamiltonian and two constants a, b, $a - b = \Omega\left(\frac{1}{poly(n)}\right)$. We are asked to decide whether the lowest eigenvalue of H is at least b or at most a, and we are promised that one of the two occurs. In the commuting case, we are guaranteed that $[H_i, H_j] = 0$ for all i, j, and W.L.O.G, we also assume that H_i's are projections. a and b can be taken to be 0 and 1 in this case, respectively.*

2.1 Interaction graphs and their expansion

One can consider various definitions for the interactions graphs underlying a local Hamiltonian.

Definition 2 ***Bi-Partite Interaction Graph*** *For a (k,d)-local Hamiltonian $H = \{H_i\}_i$ we define the bi-partite interaction graph $G = (L, R; E)$ as follows: R, the nodes on the right, correspond to the n particles of $\mathcal{H} = \mathbf{C}^{d^{\otimes n}}$ and L corresponds to the set of local terms $\{H_i\}_i$. An edge exists between a constraint $l \in L$ and a particle $r \in R$ if H_l acts non-trivially (namely, not as the identity) on r. Note that the left degree is equal to k; We denote the right degree to be D_R.*

Definition 3 ***Bi-partite small-set expansion*** *A bi-partite graph $G = (L, R; E)$ is said to be ε-small-set-expanding, if for every subset of particles $S \subseteq R$ of size $|S| \leq k$, $|\Gamma(S)| \geq |S| D_R(1 - \varepsilon)$.*

Bi-partite expanders have been defined and used e.g., in [52], [24] to construct locally-testable classical codes. Here we require expansion to hold only for small sets.

2.2 Commuting terms and C^*-algebras

We now state the lemma of Bravyi and Vyalyi [18] precisely:

Lemma 1 *Let H_i, H_j be two local terms on Hilbert space $\mathcal{H}$, $[H_i, H_j] = 0$. Let $\mathcal{H}_{int}$ denote the intersection of $supp(H_i) \cap supp(H_j)$, where $supp(H)$ is the subset of qudits examined non-trivially by H. Then, there exists a direct-sum decomposition*

$$\mathcal{H}_{int} = \bigoplus_{\alpha} \mathcal{H}_{int}^{\alpha},$$

where for each α we have:

$$\mathcal{H}_{int}^{\alpha} = \mathcal{H}_{int}^{\alpha,i} \otimes \mathcal{H}_{int}^{\alpha,j},$$

such that both H_i, H_j preserve all subspaces $\mathcal{H}_{int}^{\alpha}$, and moreover, for each α $H_i|_\alpha$ is non-trivial only on the Hilbert space $\mathcal{H}_{int}^{\alpha,i}$, whereas $H_j|_\alpha$ is non-trivial only on the Hilbert space $\mathcal{H}_{int}^{\alpha,j}$.

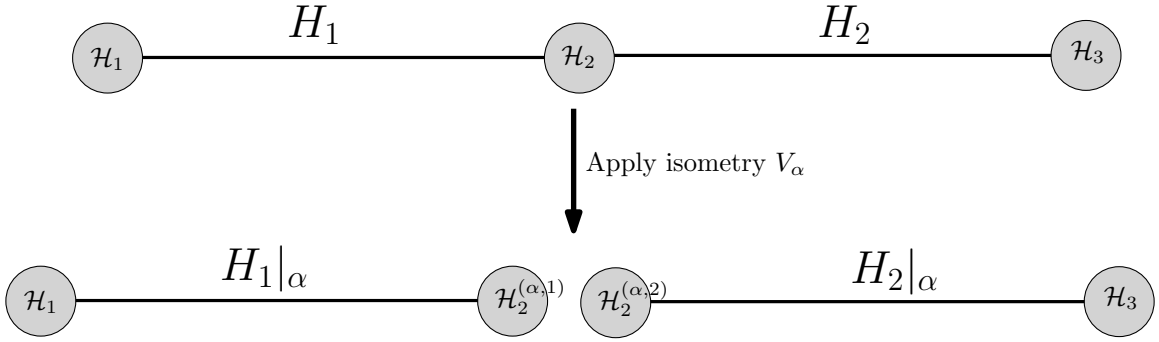

Figure 1: An example of lemma (1): a pair of 2-local Hamiltonians, are restricted to a subspace α on their intersection. Inside this subspace, they act on separate subsystems.

In this paper we apply Lemma (1) in the following context:

Corollary 1 *Let H be a $CLH(k, d)$ instance, and let H_0 be a local term such that for any $H_i \in H \setminus \{H_0\}$ that intersects H_0, the intersection is at most on one qudit. Then there exists a direct-sum decomposition, on each of the qudits of H_0, preserved by all local terms of H, such that the restriction of H_0, to any such tensor-product subspace, is disjoint from any other term in H.*

Proof: For each particle q examined by H_0, let the term H_1 be the sum of all local terms acting on q except H_0, and apply the lemma one by one on each qudit this way. ■

Note, that the above corollary does not require that any two local terms acting on a qudit in H_0, intersect only on that qudit.

2.3 stabilizer quantum error correcting codes

Definition 4 ***stabilizer Code*** *The group Π^n is the n-fold tensor product of Pauli operators $A_1 \otimes A_2 \otimes \ldots \otimes A_n$, where $A_i \in \{I, X, Y, Z\}$. along with multiplicative factors $\pm 1, \pm i$ with matrix multiplication as group operation. A stabilizer code C is defined by a set of commuting elements in Π^n; the set is denoted by $\mathcal{G}$. The group generated by $\mathcal{G}$ is an Abelian group denoted $A \subset \Pi^n$. The codespace is defined as the mutual 1-eigenspace of all elements in A (alternatively, in $\mathcal{G}$). An element of $\mathcal{E} \in \Pi^n$ is said to be an error if it does not commute*

with at least one element of $\mathcal{G}$. i.e. $\mathcal{E} \notin \mathbf{Z}(A)$, where $\mathbf{Z}(A)$ is the centralizer of A. An element of $\mathcal{E} \in \Pi^n$ is said to be a logical operation, if it commutes with all of $\mathcal{G}$, but not generated by $\mathcal{G}$, i.e., $\mathcal{E} \in \mathbf{Z}(A) - A$. A stabilizer code is said to be k-local if each term $g \in \mathcal{G}$ is a tensor product of n Pauli operators in which exactly k are different than I.

The Pauli group can be generalized to particles of any dimensionality, and thus the above definition can be generalized to work over any dimension d:

Definition 5 ***The Pauli group generalized to*** *Z_d Let $X_d^k : |i\rangle \mapsto |i + k \ modulo \ d\rangle$, $P_d^\ell |j\rangle \mapsto w_d^{j\ell} |j\rangle$ be the generalized bit and phase flip operators on the d-dimensional Hilbert space, where w_d is the primitive dth root of unity. Let Π_d be the group generated by those operators (and the necessary multiplicative factors, namely, all roots of unity of order d.) The group Π_d^n is the n-fold tensor product of Pauli operators $A_1 \otimes A_2 \otimes \ldots \otimes A_n$, where $A_i \in \left\{ X_d^k P_d^\ell \right\}$ along with the multiplicative factors.*

The definitions 4 can be generalized to any dimensionality of particles d. We can now define the distance of the code:

Definition 6 ***Distance*** *Let C be a k-local stabilizer code with generating set $\mathcal{G} \subset \Pi^n$. For an error $\mathcal{E} \in \Pi^n$, let $wt(\mathcal{E})$ denote the size of the non-trivial support of $\mathcal{E}$, and let $wt(\mathcal{E}_\mathcal{G})$ denote its minimal weight modulo the group A generated by $\mathcal{G}$, C is said to have a constant relative distance $\delta > 0$, if for any $\mathcal{E} \in \mathbf{Z}(A) - A$, we have $wt(\mathcal{E}_\mathcal{G}) \geq \delta n$.*

We note that a code with distance 1 cannot correct any error (see [34]) and so we can assume the absolute distance is strictly larger than 1. We also note that we can assume that there is no qudit q and a state on it $|\alpha\rangle$such that all states in the code look like $|\alpha\rangle$ tensor with some state on the remaining qudits; We say that in this case, the qudit is *trivial* for the code.

2.4 Notation

Throughout the paper we shall use the following notation: d is the dimension of the qudits involved. For a bi-partite graph we denote $G = (L, R; E)$ where L denotes the left set of vertices of size $|L| = m$ (corresponding to constraints in the text), R denotes the right vertices $|R| = n$ (corresponding to particles), and E is the set of edges between L and R. D will denote the degree of a graph in more generality; D_R will denote the right degree of a Bi-partite graph, which is assumed in this paper to be constant.

Given $S \subseteq R(L)$ in a bi-partite graph, $\Gamma(S)$ denotes the neighbor set of S in $L(R)$. $\mathcal{N}(q)$ will denote the neighborhood of q in R, namely all the qudits participating in all the constraints acting on q. ε (and sometimes μ) will be used to denote the expansion error for bi-partite graphs (as in Definition 3). δ will be used to denote the relative distance of a code. γ will be used to denote the promise gap, or alternatively the approximation error of a CLH instance, i.e. the fraction of terms we are allowed to throw away.

3 Approximate CLH

3.1 Approximating CLH on Expanders

We can now prove Theorem (1). As mentioned in the introduction, the proof relies on the simple important observation regarding highly expanding small set bi-partite expanders, namely, that not

too many terms intersect a term in more than one node (Claim 1). We start by proving the following:

Proof: (Of Claim 1) By definition, and by the requirement on expansion, the average degree of a vertex in $\Gamma(S)$ w.r.t. S is $\frac{D_R|S|}{|\Gamma(S)|} \leq \frac{1}{1-\varepsilon} \leq 1+2\varepsilon$, where the second inequality follows from $\varepsilon < \frac{1}{2}$. Since each vertex in $\Gamma(S)$ has degree w.r.t. S at least 1, the claim follows. ■

We now define a notion of isolation, which will help us handle single constraints:

Definition 7 ***Isolated constraints and particles*** *Let$(L, R; E)$ be a bi-partite graph. A constraint $g \in L$ is **isolated** if for any constraint $v \in L$ except g, we have $|\Gamma(v) \cap \Gamma g| \leq 1$. We define the **isolation penalty** of g, to be the minimal number of constraints in L that we need to remove so that g is isolated.*

We are now ready to prove the theorem:

Proof: (Of Theorem 1) Let H be an instance of $CLH(k, d)$, and let $(L, R; E)$ be its bi-partite interaction graph, where the degree of each $v \in R$ is D_R. We perform an iterative process. At step t, we have a set of remaining terms $L_{rem}(t)$ acting on the remaining Hilbert space $\mathcal{H}_{rem}(t)$. We also have a set $L_{bad}(t)$ which are terms we have collected so far that we want to throw away. W.l.o.g., we shall find a state whose energy approximates the ground energy of H, but this can be applied to the approximation of any eigenspace. Repeat the following:

1. **isolate** Pick an arbitrary local term v in $L_{rem}(t)$. and isolate it (Definition 7) by removing as few as possible terms from L (putting them in a set of terms which we call $L_{bad}(t)$).
2. **isometries** Now that the term is isolated, the conditions of Corollary 1 hold: By the corollary, there is a tensor product of 1-local isometries on the qudits of v: $W_v = \bigotimes_i W_i$ that is contained in the zero eigenspace of $L_{rem}(t)$, such that after applying these isometries, the term v acts on a disjoint subsystem than the rest of $L_{rem}(t)$. We conjugate v, and each term of $L_{rem}(t)$ by W_v. We then remove any triviality we encounter, be it removing qudits from Hamiltonian terms or entire terms altogether. We call this **Pruning**.
3. **update** We set $L_{rem}(t+1)$ as the set of remaining local terms of $L_{rem}(t)$ after this pruning, and then update $\mathcal{H}_{rem}(t+1)$ as the support of $L_{rem}(t+1)$. We set $L_{good}(t)$ to contain $v(t)$ restricted to the subspaces of the qudits, as in Corollary 1.

We terminate when there is no longer any intersection between two different terms of $L_{rem}(t)$. Clearly, this ends after at most polynomially many iterations, since the number of terms in $L_{rem}(t)$ decreases by at least 1 each iteration. Let the number of iterations be T. We claim that a state *approximating* the ground energy of the original Hamiltonian to within $|L_{bad}(T)|$, can be recovered from this procedure by finding the ground state of all terms in $L_{good}(T)$, and applying the inverse of the isometries applied along the way, which amounts to applying a constant-depth quantum circuit. This is true as long as the subspaces in the direct sum that were chosen contained the ground state of the current Hamiltonian; the NP prover can provide the indices of the subspaces so that this holds.

Bounding the approximation error The non-trivial part of the proof is to upper bound $|L_{bad}(T)|$. Let us measure the following ratio at each iteration t: the number of local terms moved to $L_{bad}(t)$ during step "Isolate", divided by the overall particle dimension (namely, sum of local dimensions of particles) which we reduced. We claim that this ratio is at most $2D_R\varepsilon$, and so

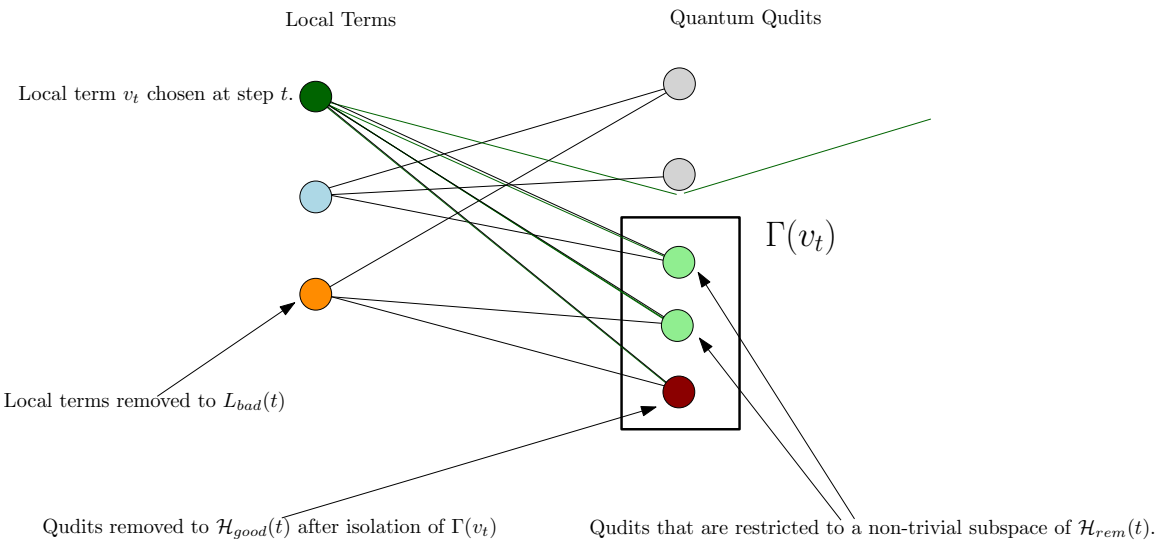

Figure 2: The isolation procedure: after removing the local terms to L_{bad}, v_t becomes isolated. We then apply local isometries to each of its qudits in $\Gamma(v_t)$. One of the qudits is exclusively examined by v_t after isolation, so it is removed altogether to $\mathcal{H}_{good}$.

Claim 2 *We remove at most $2D_R\varepsilon$ terms per local dimension we remove.*

Proof: To see this, consider a term $v(t)$ in $L_{rem}(t)$ whose support is a set S with a vertices. Let $G(t)$ denote the bi-partite interaction graph of $L_{rem}(t)$. We know that in $G(0)$, it was true that $|\Gamma(S)| \geq aD_R(1-\varepsilon)$, and so in order to isolate S, it sufficed to remove $2\varepsilon D_R a$ vertices. Since $G(t)$ is derived from $G(0)$ by removing vertices and edges, this bound still holds. So at each step, when isolating a term of size a, we remove at most $2\varepsilon D_R a$ constraints.

On the other hand, let us calculate how many qudit dimensions we remove at such an iteration, out of the total number of qudit dimensions (namely, the sum of dimensions over the particles). For each qudit $q \in S$, q is either removed altogether, and thus its local dimension is reduced by $dim(\mathcal{H}_q)$, or it is restricted to some non-trivial subspace and maybe further divided to two sub-particles, one of which is removed from $\mathcal{H}_{rem}(t)$; this act reduces the local dimension by at least 1. Thus, the total number of local dimensions removed by one application of "isolate" is at least $|S| = a$. ■

The initial total sum of particle dimensions is $d|R|$. By the claim, $|L_{bad}(T)| \leq 2D_R\varepsilon \cdot (d|R|)$. Substituting $D_R = k|L|/|R|$, we have that $|L_{bad}(T)| \leq 2k\varepsilon d|L|$, or $\frac{|L_{bad}(T)|}{|L|} \leq 2kd\varepsilon$, as desired. ■

3.2 NP-hardness of approximation

The proof of theorem 2 relies on the fact that the output of the PCP reductions in Dinur's proof [27] are excellent small set Bi-partite expanders. The gap can be amplified to, say, a $1/3$, and yet the expansion error ε can be made smaller than any constant, and certainly smaller than the required $\gamma(\varepsilon)$. This is an easy observation given the proof in [27] and we leave the details of this for the journal version.

4 Bound on the Robustness of stabilizer LTCs on Expanders

In this section we define stabilizer LTC codes and their robustness, and prove Theorem 3.

4.1 Definitions: stabilizer LTC codes

Definition 8 ***The weight of an error on stabilizer codes*** *Let C be a stabilizer code on n qudits, with generating set $\mathcal{G} \subset \Pi^n$. Let A be the associated stabilizer group. Let $\mathcal{E} \in \Pi^n$ be some element in the Pauli group. We denote by $wt(\mathcal{E})$ the number of locations in which $\mathcal{E}$ is non-identity. We denote by $wt_{\mathcal{G}}(\mathcal{E})$ the weight of $\mathcal{E}$ modulo the group generated by $\mathcal{G}$, i.e., the minimum over the weight of all words which are equal to $\mathcal{E}$ modulo an element in the group A generated by $\mathcal{G}$; Likewise, we denote by $wt_{\mathcal{Z}(\mathcal{G})}(\mathcal{E})$ the weight of $\mathcal{E}$ modulo $Z(A)$, the centralizer of A, the group generated by $\mathcal{G}$.*

We now define the robustness of a stabilizer code.

Definition 9 ***Robust Local stabilizer Codes*** *Let C be a stabilizer code, with generating set $\mathcal{G} \subset \Pi^n$, such that each particle is examined by D_R generators. C is said to be $r(\delta)$-robust, if for any $\mathcal{E} \in \Pi^n$ such that $wt_{Z(\mathcal{G})}(\mathcal{E}) = \delta n$, we have that $\mathcal{E}$ does not commute with at least $r(\delta) \cdot D_R \cdot wt_{Z(\mathcal{G})}(\mathcal{E})$ generators $g \in \mathcal{G}$ (we say that those generators, namely constraints, are violated). We call $r(\delta)$ the robustness of the code for errors of size δ. We sometimes refer to the robustness of a given error pattern $\mathcal{E}$ as r if the number of violated constraints for that error pattern is at least $r \cdot D_R \cdot wt_{Z(\mathcal{G})}(\mathcal{E})$.*

4.2 A useful fact about stabilizers

Definition 10 ***Restriction of stabilizers*** *For a Pauli operator A, let $A|_q$ denote the q-th component of the tensor product A, and let $A|_{-q}$ denote the tensor product of all terms except the q-th. Similarly, for a generating set $\mathcal{G}$, we denote by $\mathcal{G}|_q$ as the set $\{g|_q | g \in \mathcal{G}\}$, and similarly for $\mathcal{G}|_{-q}$.*

We now prove a useful fact: that the restrictions to a given qudit q of all the generators of a stabilizer code with distance larger than 1 cannot all commute.

Fact 3 *Let C be a stabilizer code with minimal distance at least 2. Then for any qudit q, there exist two generators $g(q), h(q)$ acting non-trivially on q such that $[g|_q, g'|_q] \neq 0$.*

Proof: Otherwise there is a qudit q such that for any $a, b \in \mathcal{G}$ we have $[a|_q, b|_q] = 0$. Let $Q = g|_q$ for some $g \in \mathcal{G}$ s.t. Q is not the identity. We have that $Q' = Q \otimes I_{-q}$, namely the tensor product with identity on the other qubits, commutes with all $g \in \mathcal{G}$, and thus is in the centralizer of $\mathcal{G}$: $Q' \in \mathcal{Z}(\mathcal{G})$. However, Q' cannot be insider the stabilizer, since otherwise q is in a constant state (the 1 eigenvector of Q) for all words in the code, and is thus trivial for the code (see Section 2.3). Hence, $Q' \in Z(A) - A$, where A is the stabilizer group and $Z(A)$ its centralizer, and so the distance of the code is 1. ■

4.3 Proof of Theorem 3

In the proof we will make use of "sparse" sets of constraints, defined as follows.

Definition 11 ***L-independent set of constraints*** *The neighborhood of a constraint is all constraints which intersect it in at least one qudit. A set of constraints $U \subseteq L$ is said to be L-independent if for any two constraints in the set, their two neighborhoods do not intersect on any qudit.*

Proof: (Of theorem 3)

Generating the error We want to construct an error with large weight modulo the centralizer that will not violate too many constraints in $\mathcal{G}$. Let C be a stabilizer code with a k-local generating set $\mathcal{G}$, such that the bi-partite interaction graph of S is a small-set bi-partite ε-expander. Let U be an L-independent set of constraints of size δn. We note that since $\delta \leq \frac{1}{k^3 D_R}$ an L-independent set of this size must exist, by a simple greedy algorithm. For a given $u \in U$, and $i \in [k]$, let $\alpha_i(u)$ denote the number of check terms incident on qudit i of $\Gamma(u)$ that have at least two representatives in $\Gamma(u)$. Then for each u we define $q(u)$ to be the qudit of minimal $\alpha_i(u)$ over $[k]$. Let $T = \{q(u) | u \in U\}$. Let us define an error pattern:

$$\mathcal{E} = \bigotimes_{u \in U} u|_{q(u)}.$$

We first note that $\mathcal{E}$ is not inside the centralizer $C(\mathcal{G})$, and therefore it is an error; This is true since by fact (3) for each qudit q in the support of $\mathcal{E}$, $\mathcal{E}|_q$ does not commute with $h|_q$ for some $h \in \mathcal{G}$. But since T is induced by an L-independent set, h does not touch any other qudit in the support of $\mathcal{E}$ except q, so this implies $[h, \mathcal{E}] = [h|_q, \mathcal{E}|_q] \neq 0$. We will now show that $\mathcal{E}$ has large weight modulo the group, but is penalized by a relatively small fraction of the check terms.

Weight Analysis By definition, we have that $wt(\mathcal{E}) = |T| = |U| = \delta n$. We need to show

$$wt_{C(\mathcal{G})}(\mathcal{E}) = |T| \tag{2}$$

Since δ was chosen to be smaller than half the distance of the code C, $wt_{Z(\mathcal{G})}(\mathcal{E}) = wt_{\mathcal{G}}(\mathcal{E})$ and so it suffices to argue about the weight of $\mathcal{E}$ modulo $\mathcal{G}$.

Suppose on the negative that $wt_{\mathcal{G}}(\mathcal{E}) < |T|$. Let $\mathcal{E}'$ be an error pattern which is equal to $\mathcal{E}$ modulo $\mathcal{G}$ but has weight $< |T|$. In other words, let $\Delta \in S$ be a word in the group S s.t. $\Delta\mathcal{E} = \mathcal{E}'$. Since the weight of $\mathcal{E}'$ is strictly smaller than that of $\mathcal{E}$, there must be one qudit q_0 in T, s.t. on the neighborhood $\mathcal{N}(q_0)$ the weight of $\mathcal{E}'$ is strictly smaller than that of $\mathcal{E}$, which is 1; namely, $\mathcal{E}'$ must be equal to the identity on all the neighborhood of q_0. Here, we have used here the fact that the neighborhoods of different qudits in an $L-$independent set are non-intersecting (definition 11). This means that Δ must be equal to the inverse of $\mathcal{E}$ on this neighborhood. But this inverse is exactly the following: It is equal to $\mathcal{E}|_{q_0}^{-1}$ on q_0, and to the identity on all other qudits in the neighborhood; and moreover, we have by construction that $\mathcal{E}|_{q_0}^{-1}$ on q_0, (and therefore also that $\mathcal{E}|_{q_0}$) does not commute with $h|_{q_0}$, for some $h \in \mathcal{G}$.) This means that Δ does not commute with $h \in \mathcal{G}$, in contradiction to the fact that $\Delta \in S$.

Robustness Analysis We upper-bound the number of generators that do not commute with $\mathcal{E}$. For each $u \in U$, we have that the number of generators $g \in \mathcal{G}$ that do not commute with $\mathcal{E}|_{q(u)}$, is at most the number of generators that share at least two qudits with u. By fact (2) there exists a qudit such that the fraction of its check terms with at least two qudits in $supp(u)$ is at most 2ε. Since we chose $q(u)$ to be the qudit that minimizes that fraction over all qudits of u, the absolute number of such check terms is at most $2\varepsilon D_R$. Hence the overall penalty on $\mathcal{E}$ is at most $2\varepsilon|T|D_R$. By Equation 2 we get that the penalty is at most $2\varepsilon D_R wt(\mathcal{E}_{\mathcal{G}})$. By definition 9, this implies that $r(\delta) = r(wt(\mathcal{E}_{\mathcal{G}})/n) \leq 2\varepsilon$. ■

Corollary 2 *Let C be a stabilizer code of minimal distance > 1 and a k-local succinct generating set, such that each qubit is examined by D_R constraints. If there exists an L-independent set U, s.t. $|U| = \delta n$ for some $\delta > 0$, and U's neighbor set of qudits has expansion error at most ε (i.e. $|\Gamma(\Gamma(U))| \geq |\Gamma(U)|D_R(1-\varepsilon)$), then for any $\delta' \leq \delta$ we have that $r(\delta') \leq 2\varepsilon$.*

Proof: The proof follows exactly the proof of Theorem 3 noticing that the only thing we need from the assumption on the expansion of the graph, is the existence of an L-independent set whose expansion errors is at most ε. ∎

5 An upper-bound on robustness

We now show an absolute constant, upper-bounding the robustness of any quantum stabilizer code C with local generators. We start with an easy alphabet based upper bound.

5.1 Alphabet-based bound on robustness

In attempting to understand robustness of a stabilizer code, one must first account for limitations on robustness that seem almost trivial, and occur even when there is just a single error:

Definition 12 Single error robustness *Let $f(d) = d^2$ be the number of distinct generalized Pauli's on a Hilbert space of dimension d. Let $t(d) = 1/(f(d) - 1)$; The single error robustness in dimension d is defined to be $\alpha(d) = 1 - t(d)$.*

The motivation for the above definition is as follows. For any qudit q, there always exists Q, a non-identity Pauli operator among the $d^2 - 1$ Pauli's on a d-dimensional qudit, such that a fraction at least $t(d)$ of the generators touching q are equal to Q when restricted to q. If we consider a one qudit error on q to be equal to Q, then it would commute with $t(d)$ of the generators touching q; and can at most violate $\alpha(d)$ of the constraints touching it. Thus, one can expect that it is possible to construct an error of linear weight, whose robustness is bounded by the single error robustness, using qubits whose neighborhood sets of constraints are far from each other. Indeed, we show:

Fact 4 ***Alphabet bound on robustness***
For any stabilizer code C on n d-dimensional qudits, of distance at least 2, and a k-local succinct generating set $\mathcal{G}$, whose right-degree is D_R, we have $r(\delta) \leq \alpha(d)$, for any $\delta \leq 1/(k^3 D_R)$.

Proof: Similar to theorem (3) there exists an L-independent set U of size δn. For each $u \in U$ we select some qubit $q \in \Gamma(u)$ and examine the stabilizer component on q of all stabilizers on q. Consider $P(q)$, the restriction of all generators incident on q to the qudit q, and let $MAJ(q)$ denote the Pauli that appears a maximal number of times in $P(q)$. We then set $\mathcal{E} = \bigotimes_{u \in U} MAJ(q)$. We first realize that $\mathcal{E}$ is an error: we want to show that there exists a generator g such that $\mathcal{E}$ and g do not commute. Otherwise, $\mathcal{E}$ commutes with all generators. This implies that for all q, $P(q)$ consists of a single Pauli, contradictory to fact (3). Similar to equation (2) of theorem (3) the weight of $\mathcal{E}$ is $|U| = \delta n$. Furthermore, for each qudit q, the fraction of generators on q, whose restriction to q does not commute with $\mathcal{E}|_q$ is at most $\alpha(d)$, since the number of appearances of $\mathcal{E}|_q = MAJ(q)$ in $P(q)$ is at least $t = 1 - \alpha$. Hence the number of violated constraints is at most $\alpha(d) \cdot |U| \cdot D_R$. ∎

5.2 Separation from alphabet-based robustness

In this section we show that this alphabet-based bound in fact cannot be achieved, and the robustness is further bounded and attenuated by a constant factor below the single qudit robustness, due to what seems to be a strictly quantum phenomenon. We will use the topology of the underlying graph to

achieve this separation, by treating differently expanding instances and non-expanding instances. Before stating the main theorem of this section, we require a new definition, and a simple fact:

Definition 13 ***k-independent set of constraints*** *Let H be a k-local Hamiltonian instance, and let u be a constraint of H. The 0-th neighborhood of u, denoted $\Gamma^{(0)}(u) = \Gamma(u)$ is the set of qudits examined by u. We define recursively, the t-th neighborhood of u, $\Gamma^t(u)$ to be the set of qudits, belonging to all constraints which act on qudits in $\Gamma^{(t-1)}(u)$. A set of constraints is said to be k-independent if for any $a, b \in U$ we have $\Gamma^{(k)}(u) \cap \Gamma^{(k)}(v) = \Phi$.*

The following fact can be easily derived by a greedy algorithm:

Fact 5 *Let $\eta(k, D_R) = k^{-(2k+1)} D_R^{-(2k-1)}$. For any k-local Hamiltonian H on n qudits, such that each qudit is examined by D_R local terms, there exists a k-independent set of size at least ηn.*

Proof: Pick a constraint, remove all constraints in its $\Gamma^{(2k)}$ neighborhood, and repeat. We get that the fraction of constraints is at least $k^{-(2k)} \cdot D_R^{-(2k)}$, and by the fact that the number of constraints times k, is equal to nD_R, we get the desired result. ∎

Theorem (4) *Let C be a stabilizer code on n d-dimensional qudits, of minimal distance at least k, and a k-local ($k \geq 4$) succinct generating set $\mathcal{G} \subset \Pi^n$, where the right degree of the interaction graph of $\mathcal{G}$ is D_R. Then there exists a function $\gamma_{gap} = \gamma_{gap}(k) > min\left\{10^{-3}, 0.01/k\right\}$ such that for any $\delta \leq \min\{dist(C)/2, \eta/10\}$, we have $r(\delta') \leq \alpha(d)\,(1 - \gamma_{gap})$. where $\delta' \in (0.99\delta, 1.01\delta)$.*

The proof of the theorem will use, on one hand, corollary (2) which upper-bounds the robustness of expanding instances, and on the other hand a lemma on non-expanding instances,which essentially, tries to "mimic" the behavior of the classical setting, in which non-expanding topologies suffer from poor robustness. We now state this lemma:

Lemma 2 *Let C be a stabilizer code on n qudits of dimension d, with minimal distance at least k and a k-local ($k \geq 4$) succinct generating set $\mathcal{G}$, where the right degree of the interaction graph of $\mathcal{G}$ is D_R. Let $\gamma_{gap} = \gamma_{gap}(k) = min\left\{10^{-3}, 0.01/k\right\}$. If there exists a k-independent set U of size $|U| = \delta n$, with $\delta < \min\{dist(C)/2\}$, such that the bi-partite expansion error of $\Gamma(U)$ is at least $\varepsilon = 0.32$, i.e. $|\Gamma(\Gamma(U))| = |\Gamma(U)|D_R(1 - \varepsilon')$ for some $\varepsilon' \geq 0.32$ then*

$$r(\delta') \leq \alpha(d) \cdot (1 - \gamma_{gap}),$$

for some $\delta' \in 0.1(0.99\delta, 1.01\delta)$.

From this lemma, it is easy to show theorem (4):

Proof: (of theorem 4) The parameters of the theorem allow us to apply directly fact (5); hence there exists a k-independent set S of size at least ηn, for η as defined in fact(5). Hence, since $\delta \leq \eta/10$ there exists a k-independent set S of size 10δ. Now, either:

1. S has expansion error at least 0.32. By lemma (2), for any $k \geq 4$ we have

   $$r(\mu) < \alpha(d)(1 - \gamma_{gap}),$$

   for some $\mu \in 0.1(0.99 \cdot 10\delta, 1.01 \cdot 10\delta) = (0.99\delta, 1.01\delta)$, and $\gamma_{gap}(k)$ from lemma (2), which is at least $min\left\{10^{-3}, 0.01/k\right\}$.

2. The set S is ε-expanding for $\varepsilon < 0.32$. In which case, since S is in particular L-independent, then by corollary (2), the robustness $r(\delta') \leq 2\varepsilon < 2/3 - 0.01 \leq \alpha(d) - 0.01$, for all $\delta' \leq |S|/n$. In particular $r(\mu) < \alpha(1 - 0.01/k)$.

Taking the higher of these two bounds we get the desired upper-bound for $r(\mu)$. ■

5.3 Proof of the lemma (2)

5.3.1 General structure of proof

Let us clarify what we're trying to show. We want to show that if the expansion is bad, errors cannot have large relative penalties. Consider a set S with positive expansion error $\varepsilon > 0$. A-priori, if we have an error on S, then the maximal number of violations is strictly less than $|S|D_R$, and in fact at most $|S|D_R(1-\varepsilon)$. This might seem as though it proves the lemma trivially.

The technical problem here, however, is that an error on S may just "seem" to be large, whereas possibly, may be represented much more succinctly modulo the stabilizer group. We would hence like to devise an error pattern, that cannot be downsized significantly, but would still "sense" the non-expanding nature of S, and hence have fewer-than-optimal violations. We start with a fact lower-bounding the weight modulo the group of an error confined to the qudits of a single generator; we call this the Onion fact since its proof (given in Subsection 5.3.4) works via some hybrid argument on the onion-like layers $\Gamma^{(i)}(u)$ surrounding the qudits of one term u.

Fact 6 ***Onion fact***
Let C be a stabilizer code on n qudits with a succinct generating set $\mathcal{G}$ of locality k, such that $dist(C) \geq k$. Let $\mathcal{E}$ be some word in Π^n s.t. $supp(\mathcal{E}) \subseteq \Gamma^{(0)}(u) = \Gamma(u)$ for some generator $u \in \mathcal{G}$. Finally let $\Delta \in C$ be some word in the code and let $\mathcal{E}_\mathcal{G} = \Delta\mathcal{E}$. Then, for any $i \in [k]$, if $wt(\mathcal{E}|_{\Gamma(u)}) = i$, then $wt(\mathcal{E}_\mathcal{G}|_{\Gamma^{(k)}(u)}) \geq min\{i, k-i\}$.

Now, let us see what the Onion fact means. It states that given an error on the k-qudit support of a generator, the weight of any representation of this error modulo the centralizer, cannot be reduced in the k-neighborhood of the generator, provided that the error has weight at most $k/2$.

Our idea is to concentrate the error on a large set of such "islands", each island supported on one generator, s.t. the generators supporting those islands are sufficiently far away from each other (in the interaction graph) so that the k-neighborhoods of those generators are non-intersecting.

If we draw a random error on the qudits on these "islands", such that the expected number of errors per "island", is say 1 error, the following will occur: on one hand, there will be a good portion of "islands" with at least two errors, and these two errors will "sense" the sub-optimality of number of neighbors due to the expansion, interfere with each other, and cause an overall reduced penalty of that "island". On the other hand, only a meager fraction, exponentially small in k, of those "islands" with at least two errors, will have more than $k/2$ errors; only those, by the Onion fact (fact 6) can reduce their weight modulo the centralizer.

We calibrate our parameters in the random choice of our error pattern to have expected number of errors which is indeed order of 1, so that the above tradeoff will indeed hold.

In the following we first define the error; We provide the proof that the expected penalty of this error is small in fact (7), then prove the onion fact in Subsubsection 5.3.4 and using it we then prove Fact (8), in which we show that the error has large weight modulo the group. Finally we combine all the above to finish the proof of the lemma.

5.3.2 Constructing the error

Let $U \subseteq L$ be a k-independent set as promised by the conditions of the lemma. Then $|U| = \delta n$, and denoting $S = \Gamma(U)$, we have that $|S| = \delta nk$. Therefore, $|\Gamma(S)| = |S|D_R(1-\varepsilon')$, for some $\varepsilon' \geq 0.32$. Let $\mathcal{E}$ be the following random error process: for each qudit of S independently, we apply I w.p. $1-p$ for $p = 1/(10k)$, and one of the other Pauli operators: with equal probability $p \cdot t(d)$, where t is defined in (12).

$$\mathcal{E} = \bigotimes_{i \in S} \mathcal{E}_i, \text{ where } \mathcal{E}_i = \begin{cases} I_i & \text{w.p. } 1 - 1/(10k) \\ X_d^k P_d^l & \text{w.p. } t/(10k) \end{cases}$$

We note here that the choice of p is such that on average, each k-tuple has only a small number of errors; the expectation of the number of errors is an absolute constant $1/10$ (not a fraction of k). This will help, later on, to lower-bound the weight of the error modulo the group.

5.3.3 Analyzing Penalty

We first claim, that on average, $\mathcal{E}$ has a relatively small penalty w.r.t. $\mathcal{G}$, using the fact that the expansion error is at least 0.32 as in the condition of Lemma 2

Fact 7

$$\mathbf{E}_{\mathcal{E}}\left[Penalty\right] \leq p\alpha|S|D_R\left(1 - 0.02/k\right)$$

Proof: Let $G = (L, R; E)$ denote the bi-partite graph corresponding to $\mathcal{G}$, with L being the constraints and R the qudits. Let $S = \Gamma(U)$ be as before. Let the error process $\mathcal{E}$ be the one defined above. For any constraint $c \in \Gamma(S)$ which is violated when applied to this error, observe that there must be a qudit $i \in supp(c)$ such that $[c|_i, \mathcal{E}_i] \neq 0$. We now would like to bound the number of constraints violated by $\mathcal{E}$ using this observation, and linearity of expectation.

For an edge $e \in E$ connecting a qudit i in S and a constraint c in $\Gamma(S)$, let $x(e)$ denote the binary variable which is 1, iff the error term $\mathcal{E}_i$ on does not commute with $c|_i$. In other words, an edge marked by 1 is an edge whose qudit may cause its constraint to be violated. By construction, for each $e \in E$ which connects the qudit i and the constraint c we have

$$\mathbf{E}_{\mathcal{E}}[x(e)] = p(1-t). \tag{3}$$

This is true since a constraint c restricted to the qudit i, $c|_i$ does not commute with the error restricted to the same qudit i, $\mathcal{E}_i$, iff both $\mathcal{E}_i$ is non-identity (which happens with probability p) and is not equal to $c|_i$.

If we had just added now $x(e)$ over all edges going out of S (whose number is $|S|D_R$), then by linearity of expectation, this would have given an upper bound on the expected number of violated constraint equal to

$$\sum_e p(1-t) = p|S|D_R\alpha(d). \tag{4}$$

Unfortunately this upper bound does not suffice; to strengthen it we would now like to take advantage of the fact that many of those edges go to the same constraint, due to the fact that the expansion is bad; thus, instead of simply summing these expectation values, we take advantage of

the fact that two qudits touching the same constraint cannot contribute twice to its violation. We note that the bound we get does not gain back the full factor of $1-\varepsilon$ but a worse one.

Let $E_{inj} \subseteq E$ be a subset of the edges between S to $\Gamma(S)$ chosen by picking one edge for each constraint in $\Gamma(S)$. For an edge $e \in E$ let $c(e)$ denote the constraint incident on e, and let $e_{inj}(c(e))$ denote the edge in E_{inj} that is connected to $c(e)$.

We now bound the expectation by subtracting $x(e)$ from the sum, if the Boolean variable $x(e_{inj}(c(e)))$ is 1; this avoids counting the violation of the same constraint twice due to the two edges. We have:

$$\mathbf{E}_{\mathcal{E}}\left[Penalty\right] \leq \mathbf{E}_{\mathcal{E}}\left[\sum_{e\in E_{inj}} x(e) + \sum_{e\notin E_{inj}} (1 - x(e_{inj}(c(e)))) \cdot x(e)\right].$$

Note that it may even be the case that some edges may cause constraints to become "unviolated", so the actual bound may be even lower. Expanding the above by linearity of expectation:

$$\mathbf{E}\left[Penalty\right] \leq \sum_{e\in E_{inj}} \mathbf{E}_{\mathcal{E}}\left[x(e)\right] + \sum_{e\notin E_{inj}} \mathbf{E}_{\mathcal{E}}\left[x(e)\right] - \sum_{e\notin E_{inj}} \mathbf{E}_{\mathcal{E}}\left[x(e_{inj}(c(e))) \cdot x(e)\right] =$$

$$\sum_{e\in E} \mathbf{E}_{\mathcal{E}}\left[x(e)\right] + \sum_{e\notin E_{inj}} \mathbf{E}_{\mathcal{E}}\left[x(e_{inj}(c(e))) \cdot x(e)\right].$$

We have already calculated the first term in the sum in Equation 4; We now lower bound the correction given by the second term. We use the fact that for any $e \notin E_{inj}$

$$\mathbf{E}_{\mathcal{E}}\left[x(e_{inj}(c(e))x(e)\right] = \mathbf{E}_{\mathcal{E}}[x(e_{inj}(c(e)))]\mathbf{E}_{\mathcal{E}}[x(e)]$$

since $\mathcal{E}$ is independent between different qudits. We can thus substitute Equation 3, and get:

$$\mathbf{E}_{\mathcal{E}}\left[Penalty\right] \leq p\alpha|S|D_R - |S|D_R\varepsilon(p\alpha)^2.$$

where we have used the fact that $|E_{inj}| = |S|D_R\varepsilon$. This is equal to

$$p\alpha|S|D_R(1 - p\alpha\varepsilon).$$

Using $p = 1/(10k), \varepsilon \geq 0.32$, $\alpha(d) \geq 2/3$, we get the desired bound. ∎

5.3.4 Towards analyzing the weight: Proof of the Onion fact

Proof: (Of Fact 6) If $\Delta|_{\Gamma(u)} = I$ then

$$wt\left(\mathcal{E}_{\mathcal{G}}|_{\Gamma^{(k)}(u)}\right) = wt\left(\mathcal{E}_{\mathcal{G}}|_{\Gamma(u)}\right) = wt\left(\mathcal{E}|_{\Gamma(u)}\right) = i. \tag{5}$$

Otherwise, $\Delta|_{\Gamma(u)}$ is non-identity, and so has at least one non-identity coordinate, and also, since Δ is non-identity, by the assumption on the succinctness of $\mathcal{G}$ we have $wt(\Delta) \geq k$.

Moreover, we claim that $wt\left(\Delta|_{\Gamma^{(k)}(u)}\right) \geq k$. Otherwise, consider the following process. Start with the constraint u, and consider the qudits in $\Gamma(u) = \Gamma^{(0)}(u)$. Now add the qudits belonging to all constraints in $\Gamma^{(1)}(u)$; Then add the next level, and so on until we have added add qudits belonging to $\Gamma^{(k)}(u)$. By the pigeonhole principle, if $wt\left(\Delta|_{\Gamma^{(k)}(u)}\right) < k$, then there must exist a level t s.t. $1 \leq t \leq k$ where Δ has zero support on qudits added in this level. This is in contradiction to the

fact that the distance of the code is at least k, since we claim that $\tilde{\Delta} = \Delta|_{\Gamma^{(t-1)}(u)}$, is in the centralizer $\mathbf{Z}(\mathcal{G})$ but its weight is less than k. To see that $\tilde{\Delta}$ is in the centralizer, we observe first that Δ commutes with all elements of $\mathcal{G}$ that act only on qudits in $\Gamma^{(t-1)}(u)$, and since $\tilde{\Delta}$ agrees with Δ on $\Gamma^{(t-1)}(u)$, $\tilde{\Delta}$ also commutes with them. We also observe that $\tilde{\Delta}$ trivially commutes with all elements in $\mathcal{G}$ whose support does not intersect $\Gamma^{(t-1)}(u)$. Hence we only need to worry about those terms that act on at least one qudit in $\Gamma^{(t)}(u) - \Gamma^{(t-1)}(u)$ and at least one qudit in $\Gamma^{(t-1)}(u)$. Let v be some such term. Note that v cannot act on any qudit outside $\Gamma^{(t)}(u)$ by definition (of the $\Gamma^{(i)}(u)$'s). We know that Δ commutes with v. But by the choice of t, we know that Δ is trivial on those qudits added at the tth level, and hence Δ restricted to $\Gamma^{(t)}(u)$ (which contains the qudits of v) is the same as Δ restricted to $\Gamma^{(t-1)}(u)$. And so Δ restricted to $\Gamma^{(t-1)}(u)$ commutes with v.

We showed that $\tilde{\Delta}$ is in $\mathbf{Z}(\mathcal{G})$. If it also belongs to $\mathcal{G}$, this contradicts succinctness of $\mathcal{G}$; otherwise it is in $\mathbf{Z}(\mathcal{G}) - \mathcal{G}$ implying the distance of C is at most $k-1$, contrary to assumption. This means that $wt\left(\Delta|_{\Gamma^{(k)}(u)}\right) \geq k$. Therefore, we now know that

$$wt\left(\mathcal{E}_{\mathcal{G}}|_{\Gamma^{(k)}(u)}\right) \geq wt\left(\Delta|_{\Gamma^{(k)}(u)}\right) - wt\left(\mathcal{E}|_{\Gamma^{(k)}(u)}\right) = \tag{6}$$

$$wt\left(\Delta|_{\Gamma^{(k)}(u)}\right) - wt\left(\mathcal{E}|_{\Gamma^{(0)}(u)}\right) \geq k - i.$$

Taking the minimal of the bounds from Equations (5),(6) completes the proof. ■

5.3.5 Analyzing error weight

We note that the expected weight of $\mathcal{E}$ is $p|S|$ and since $|S|$ is linear in n, by Chernoff the probability that the weight of $\mathcal{E}$ is smaller than by a constant fraction than this expectation is $2^{-\Omega(n)}$. We need to show a similar bound on the weight modulo the centralizer group; given that $\delta < dist(C)/2$ we only need to bound the weight modulo the stabilizer group. Let $\Delta \in A$ be some element in the stabilizer group and let $\mathcal{E}_{\mathcal{G}} = \Delta\mathcal{E}$. We now need to lower-bound $wt(\mathcal{E}_{\mathcal{G}})$.

Fact 8 *For integer k, let $\hat{k} = \lfloor k/2 \rfloor + 1$. Let $y(k) : [4, \infty] \mapsto \mathbf{R}$ be the function:*

$$y(k) = \begin{cases} 1 - 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+4.54} & k \geq 12 \\ 0.9999 & 6 \leq k \leq 11 \\ 0.9992 & k = 5 \\ 0.9985 & k = 4 \end{cases}$$

We claim:

$$Prob_{\mathcal{E}}\left(wt(\mathcal{E}_{\mathcal{G}}) < |S|py(k)\right) = 2^{-\Omega(n)}.$$

Proof: Let $x \sim B(k, p = 1/(10k))$ denote a random variable which is the sum of k i.i.d Boolean variables, each equal to 1 with probability p; in other words, x is a binomial process; $B(i) = Prob(x = i)$. Let $U_i = \left\{u \in U | wt(\mathcal{E}|_{\Gamma(u)}) = i\right\}$ be the set of constraints in which exactly i errors occurred. Using the Hoeffding bound, for a given $i \in [k]$ and a given constant $\chi > 0$, we have

$$Prob_{\mathcal{E}}\left(\left|\frac{|U_i|}{|U|} - B(i)\right| \geq \chi\right) = 2^{-\Omega(n)} \tag{7}$$

By the union bound, we have that for any constant $\chi > 0$:

$$Prob_{\mathcal{E}}\left(\exists i, s.t. \left|\frac{|U_i|}{|U|} - B(i)\right| \geq \chi\right) = 2^{-\Omega(n)}. \tag{8}$$

Since the set U is a k-independent set, then the sets $\left\{\Gamma^{(k)}(u)\right\}_{u\in U}$ are non-intersecting so

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \sum_{u\in U} wt\left(\mathcal{E}_{\mathcal{G}}|_{\Gamma^{(k)}(u)}\right), \tag{9}$$

By the onion fact (6), for each $u \in U_i$ we have $wt\left(\mathcal{E}_{\mathcal{G}}|_{\Gamma^{(k)}(u)}\right) \geq min\{i, k-i\}$, hence

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \sum_{i\in[k]} |U_i| min\{i, k-i\} = \frac{|S|}{k} \sum_{i\in[k]} \frac{|U_i|}{|U|} min\{i, k-i\}$$

using $k|U| = |S|$.

Therefore, using equation (8) w.p. close to 1 we have

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \frac{|S|}{k} \sum_{i\in[k]} (B(i) - \chi) min\{i, k-i\} \geq \frac{|S|}{k}\left(\sum_{i\in[k]} B(i) min\{i, k-i\} - 2^{-k^2}\right), \tag{10}$$

for $\chi = 2^{-k^2}/k^2$.

We separate the rest of the proof to two cases: $k \geq 12$ and $4 \leq k < 12$. We start with the case $k \geq 12$. Recall $\hat{k} = \lfloor k/2 \rfloor + 1$. Let

$$A_{loss} = \sum_{i\geq\hat{k}} B(i)(2i-k).$$

Then by equation (10) we have that with probability exponentially close to 1

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \frac{|S|}{k}\left(\sum_{i\in[k]} B(i)i - A_{loss} - 2^{-k^2}\right) = \frac{|S|}{k}\left(pk - 2^{-k^2} - A_{loss}\right) \tag{11}$$

In the rest of the proof for $k \geq 12$ we upper-bound A_{loss} and substitute in the above equation to derive the desired result. Using an upper-bound of the binomial, we have:

$$B(\hat{k}) = \binom{k}{\hat{k}} p^{\hat{k}}(1-p)^{\hat{k}} \leq 2^k \cdot (10k)^{-\hat{k}}(1-p)^{\hat{k}} \leq k^{-\hat{k}} 10^{-\hat{k}} 2^k \leq 2^{-\hat{k}log(k)+k-3.3\hat{k}}, \tag{12}$$

For any $i \geq \hat{k}$ and $p < 1/2$ we have

$$B(i+1) = B(i)\left(\frac{k-i}{i+1}\right)\left(\frac{p}{1-p}\right) < B(i)\frac{p}{1-p} < 2pB(i) \tag{13}$$

Substituting equations (13) and (12) in the expression for A_{loss} we have:

$$A_{loss} = \sum_{i\geq\hat{k}}^{k} B(i)(2i-k) \leq 2^{-\hat{k}log(k)+k-3.3\hat{k}} \sum_{i\geq\hat{k}}^{k} (2p)^{(i-\hat{k})}(2i-k) \tag{14}$$

$$\leq 2^{-\hat{k}log(k)+k-3.3\hat{k}+1+\hat{k}} \sum_{i\geq\hat{k}}^{k} (p)^{(i-\hat{k})}(i - \lfloor k/2 \rfloor) \tag{15}$$

Changing summation $i - \lfloor k/2 \rfloor \mapsto j$ we have the above is at most:

$$2^{-\hat{k}log(k)+k-2.3\hat{k}+1} \sum_{j\geq 1}^{\lceil k/2 \rceil} p^{-j+1} j \leq 2^{-\hat{k}log(k)+k-2.3\hat{k}+1} \sum_{j\geq 1}^{\lceil k/2 \rceil} p^{-j+1} k \tag{16}$$

$$\leq 2^{-\hat{k}log(k)+k-2.3\hat{k}+1} k \sum_{j\geq 1}^{\lceil k/2 \rceil} p^{-j+1} \leq 2^{-\hat{k}log(k)+k-2.3\hat{k}+1} k \cdot 1.1 \leq 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+1.2}, \tag{17}$$

where in the last inequality we bound the sum by $\sum_{i\geq 0} 1/p^i$, and set $p = 1/(10k) \leq 1/100$, using $k \geq 12$. Substituting this value in (11) we have that with probability $2^{-\Omega(n)}$ close to 1,

$$wt(\mathcal{E}_{\mathcal{G}}) \geq \frac{|S|}{k} \left(pk - 2^{-k^2} - 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+1.2} \right) =$$

$$\geq \frac{|S|}{k} \left(pk - 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+1.21} \right)$$

where in the last inequality we used again $k \geq 12$. Continuing, using $p = \frac{1}{10k}$ the above bound is equal to

$$= |S|p \left(1 - 2^{(-\hat{k}+1)log(k)+k-2.3\hat{k}+1.21+log_2(10)} \right) \geq |S|py(k),$$

for all $k \geq 12$. For values of $4 \leq k < 12$ we substitute directly k in equation (10), evaluate, and show it is at least $|S|py(k)$.

■

5.3.6 Concluding the proof of lemma (2)

Proof: By fact (7) the average penalty of $\mathcal{E}$ is small on average, i.e.

$$\mathbf{E}\left[Penalty(\mathcal{E})\right] \leq |S|D_R p\alpha(1 - 0.02/k) \triangleq P.$$

Yet, by fact (8) w.p. exponentially close to 1, we have

$$wt(\mathcal{E}_{\mathcal{G}}) \geq |S|py(k) \triangleq W_{low} \geq |S|p \cdot 0.99.$$

Similarly, by the Hoeffding bound w.p. exponentially close to 1, we have

$$wt(\mathcal{E}_{\mathcal{G}}) < |S|p(1 + 0.01) \triangleq W_{high}.$$

Since all penalties are non-negative, we conclude that *conditioned* on $|wt(\mathcal{E}_{\mathcal{G}})/(|S|p) - 1| < 0.01$, we have $\mathbf{E}\left[Penalty(\mathcal{E})\right] \leq P + 2^{-\Omega(n)}$. Therefore, there must exist an error $\mathcal{E}$, whose weight modulo $\mathcal{G}$ deviates by a fraction at most 0.01 from $|S|p$, and whose penalty is at most $P + 2^{-\Omega(n)}$.

We would like to bound the robustness of this error, which is the ratio of the penalty to its relative weight times D_R. We get that its robustness is at most

$$r = \frac{P + 2^{-\Omega(n)}}{D_R W_{low}} \leq \frac{1}{D_R} \cdot \frac{|S|D_R p\alpha(1 - 0.019/k)}{|S|py(k)} = \alpha \left(\frac{1 - 0.019/k}{y(k)} \right). \tag{18}$$

We now note that in the last expression, for all $k \geq 12$, the ratio $\frac{1-0.019/k}{y(k)}$ is at most $1 - 0.01/k$. For all values of $4 \leq k < 12$ we substitute the appropriate value of $y(k)$ and get similarly that the ratio $\frac{1-0.019/k}{y(k)}$ is at most $1 - 10^{-3}$. Hence, the robustness of the error, r is at most $\alpha(d)(1 - \gamma_{gap})$ where γ_{gap} is as defined in the statement of Theorem 4.

■

6 Acknowledgements

The authors would like to thank Gil Kalai for useful discussions regarding expander graphs; and Eli Ben-Sasson and Irit Dinur for discussions about PCPPs and LTCs.

References

[1] S. Aaronson. *The quantum PCP manifesto. http://www.scottaaronson.com/blog/?p=139 (2006).*

[2] M. Aguado, G. Vidal *Entanglement renormalization and topological order. Physical Review Letters. 03/2008; 100(7):070404.*

[3] D. Aharonov, L. Eldar On the complexity of Commuting Local Hamiltonians, and tight conditions for Topological Order in such systems *arXiv:1102.0770v2*

[4] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani. *The Detectability Lemma and Quantum Gap Amplification. Proceedings of the 41st annual ACM symposium on Theory of computing,(New York, NY, USA), pp. 417426, ACM, 2009.*

[5] N. Alon, J. Bruck, J. Naor, M. Naor, and R. Roth. *Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. IEEE Transactions on Information Theory, 38:509-516, 1992.*

[6] I. Arad. *A note about a partial no-go theorem for quantum PCP. Quantum Information and Computation 11, 1019 (2011).*

[7] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. *Proof verification and intractability of Approximation problems. J. ACM, 45(3):501-555, 1998.*

[8] S. Arora, S. Safra *Probabilistic checking of proofs: A new characterization of NP Journal of the ACM, 45 (1): 70122, 1998*

[9] S. Arora and B. Barak. *Computational Complexity: A Modern Approach. Cambridge University Press, 2009*

[10] S. Arora, B. Barak, D. Steurer *Subexponential Algorithms for Unique Games and Related Problems. FOCS 2010: 563-572*

[11] N. Bansal, S. Bravyi, and B.M. Terhal. *Classical approximation schemes for the ground-state energy of quantum and classical Ising spin Hamiltonians on planar graphs. Quant. Inf. Comp. Vol. 9, 0701 (2009).*

[12] E. Ben-Sasson, P. Harsha, and S. Raskhodnikova *Some 3CNF properties are hard to test STOC, 2003, pp. 345-354*

[13] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, S. P. Vadhan: *Robust PCPs of Proximity, Shorter PCPs, and Applications to Coding. SIAM J. Comput. 36(4): 889-974 (2006)*

[14] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, S. Vadhan *Robust PCPs of proximity, shorter PCPs, and applications to coding SIAM J. COMPUT. 2006, Society for Industrial and Applied Mathematics,Vol. 36, No. 4, pp. 889-974*

[15] E. Ben-Sasson, M. Sudan *Robust Locally Testable Codes and Products of Codes Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques,pp. 286-297, 2004*

[16] Blum, Luby, Rubinfeld *Linearity Testing / Testing Hadamard codes Encyclopedia of algorithms, Springer, (2008).*

[17] F. Brandão, A. Harrow *Approximation Guarantees for the Quantum Local Hamiltonian Problem and Limitations for Quantum PCPs* Pre-Print

[18] S. Bravyi, M. Vyalyi *Commutative version of the k-local Hamiltonian problem and common eigenspace problem Quantum Inf. and Comp., Vol. 5, No. 3 (2005) pp.187-215*

[19] S. Bravyi, M. B. Hastings, S. Michalakis *Topological quantum order: Stability under local perturbations J. Math. Phys. 51, 093512 (2010)*

[20] S. Bravyi, M. B. Hastings *A Short Proof of Stability of Topological Order under Local Perturbations. Communications in Mathematical Physics. November 2011, Volume 307, Issue 3, pp 609-627.*

[21] S. Bravyi, D. Poulin and B.M. Terhal, *Tradeoffs for reliable quantum information storage in 2D systems arxiv.org: 0909.5200 (2009), Phys. Rev. Lett. 104, 050503 (2010)*

[22] S. Bravyi, D. P. DiVincenzo, D. Loss, and B. M. Terhal. *Simulation of Many-Body Hamiltonians using Perturbation Theory with Bounded-Strength Interactions Phys. Rev. Lett. 101, 070503 (2008).*

[23] A. R. Calderbank, P. W. Shor *Good quantum error-correcting codes exist Phys. Rev. A, vol. 54, pp. 1098-1105, Aug. 1996.*

[24] M. R. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson. *Randomness conductors and constant-degree lossless expanders. In Proceedings of the ACM Symposium on Theory of Computing (STOC), pages 659-668, 2002.*

[25] A. Cross, G. Smith, J.A. Smolin, B. Zeng *Codeword Stabilized Quantum Codes. IEEE Transactions on Information Theory, Vol. 55, No.1, pp. 433-438, (2009)*

[26] E. Dennis, A. Kitaev, A. Landahl, J. Preskill *Topological quantum memory. J. Math. Phys. 43, 4452 (2002)*

[27] I. Dinur *The PCP theorem by gap amplification. STOC 2006: 241-250*

[28] I. Dinur, T. Kaufmann *On the structure of NP-hard 3-SAT instances , and a similar question for LTCs Talk at the The Fourth Israel CS Theory Day Thursday, March 24 th, 2011*

[29] I. Dinur, T. Kaufman *Locally Testable Codes and Expanders Pre-Print.*

[30] P. Erdős, R. Rado. *Intersection theorems for systems of sets Journal of the London Mathematical Society, 2nd. series 35 (1), (1960), 85-90.*

[31] E. Fetaya *Bounding the distance of quantum surface codes. J. Math. Phys. 53, 062202 (2012).*

[32] J. Friedman, A. Wigderson *On the second eigenvalue of hypergraphs Combinatorica, Volume 15, Issue 1 , pp. 43-65 ,1995*

[33] S. Gharibian, J. Kempe *Hardness of Approximation for Quantum Problems ICALP (1) 2012: 387-398*

[34] D. Gottesman *Stabilizer Codes and Quantum Error Correction Caltech Ph.D. Thesis*

[35] J. Haah *Local stabilizer codes in three dimensions without string logical operators Phys. Rev. A 83, 042330 (2011)*

[36] J. Håstad *Some optimal inapproximability results. Journal of ACM, 48:798859, 2001.*

[37] M. B. Hastings Trivial Low Energy States for Commuting Hamiltonians, and the Quantum PCP Conjecture *arXiv:1201.3387v1*

[38] M. B. Hastings, Freedman. *Pre-print, 2013.*

[39] S. Hoory, N. Linial, A. Wigderson Expander Graphs and Their Applications *Bulletin (New Series) Of The American Mathematical Society Volume 43, Number 4, October 2006.*

[40] G. Kalai, A. Lubotzky Personal Comm.

[41] S. Khot (2002) *On the power of unique 2-prover 1-round games Proceedings of the 17th Annual IEEE Conference on Computational Complexity, pp. 25*

[42] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and quantum computation, volume 47 of Graduate Studies in Mathematics. AMS, 2002.*

[43] A. Yu. Kitaev *Fault-tolerant quantum computation by anyons. Ann. Physics,303(1), 2003.*

[44] A. Yu. Kitaev *Ann. Phys. 321 2 (2006).*

[45] A. Yu. Kitaev *Annals Phys. 303, 2 (2003), arXiv:quant-ph/9707021.*

[46] A.A. Kovalev and L.P. Pryadko, Fault tolerance of "bad" quantum low density parity check codes, arXiv:1208.2317

[47] A.A. Kovalev and L.P. Pryadko, Quantum "hyperbicycle" low density parity check codes with finite rate; arXiv:1212.6703

[48] O. Parzanchevski, R. Rosenthal, and R. J. Tessler *Isoperimetric Inequalities in Simplicial Complexes arXiv:1207.0638*

[49] P. Raghavendra, D. Steurer *Graph expansion and the unique games conjecture Proceedings of the 42nd ACM symposium on Theory of computing, Pages 755-764 , STOC, 2010.*

[50] D. Schlingemann, R.F. Werner *Quantum error-correcting codes associated with graphs. Phys. Rev. A 65, 012308 (2001)*

[51] N. Schuch Complexity of commuting Hamiltonians on a square lattice of qubits *Quantum Inf. Comput. 11, 901 (2011)*

[52] M. Sipser, D. Spielman. *Expander codes. IEEE Transactions on Information Theory 42(6): 1710-1722 (1996)*

[53] A. M. Steane *Error correcting codes in quantum theory Phys. Rev. Lett., vol. 77, pp. 793-767, July 1996.*

[54] A. M. Steane *Multiple particle interference and quantum error correction Proc. Roy. Soc. Lond. A, vol. 452, pp. 2551-2577, Nov. 1996.*

[55] L. Trevisan *Approximation algorithms for unique games In Proceedings of the 46th IEEE Symposium on Foundations of Computer Science.*

[56] B. Yoshida *Feasibility of self-correcting quantum memory and thermal stability of topological order. Annals of Physics 326, 2566-2633 (2011)*

[57] G. Zémor. *On expander codes. IEEE Transactions on Information Theory, 47(2):835-837,2001.*

[58] G. Zémor *A construction of quantum LDPC codes from Cayley graphs. IEEE International Symposium on Information Theory Proceedings (ISIT), 2011, pp. 643-647*

[59] G. Zémor *Quantum LDPC codes with positive rate and minimum distance proportional to $n^{1/2}$. IEEE International Symposium on Information Theory (ISIT) 2009, pp. 799-803.*

7 Appendix

For $S \subseteq R$ let $\Gamma_1(S) \subseteq \Gamma(S)$ denote the subset of the neighbors of S with exactly one neighbor in S. Similarly, let $\Gamma_{\geq 2}(S)$ denote the subset of neighbors with at least two neighbors in S.

Proof of fact(1):

Proof: The average degree of a vertex in $\Gamma(S)$ w.r.t. $|S|$ is at most $\frac{|S|D_R}{|S|D_R(1-\varepsilon)} = \frac{1}{1-\varepsilon}$. Let α_1 denote the fraction $|\Gamma_1(S)|/|\Gamma(S)|$. Then

$$\frac{1}{1-\varepsilon} = \alpha_1 1 + (1-\alpha_1)m,$$

where m is the average degree of a vertex with at least two neighbors in S. Then

$$\alpha_1 = 1 - \frac{\varepsilon m}{m-1}.$$

Since $m \geq 2$, then α_1 is minimized for $m = 2$, and therefore

$$\alpha_1 \geq 1 - 2\varepsilon.$$

■

Proof of fact (2):

Proof: By definition, we have $|\Gamma(S)| \geq |S|D_R(1-\varepsilon)$. Let $E_{inj} \subseteq E(S)$ be a subset of the edges incident on S such that each $u \in \Gamma(S)$ has a single neighbor in S connected by an edge of E_{inj}. Then E_{inj} is of size $\Gamma(S)$ which is at least $|S|D_R(1-\varepsilon)$. Also $|E(S)| = |S|D_R$, thus $|E(S) - E_{inj}| \leq |S|D_R\varepsilon$. Therefore $|\Gamma_{\geq 2}(S)| \leq |S|D_R\varepsilon$. Hence, $\Gamma_1(S) = \Gamma(S) - \Gamma_{\geq 2}(S)$ is of size at least $|S|D_R(1-\varepsilon) - |S|D_R\varepsilon = |S|D_R(1-2\varepsilon)$. Therefore, there exists a vertex $v \in S$ with at least $D_R(1-2\varepsilon)$ neighbors in $\Gamma_1(S)$. Since v has D_R neighbors in $\Gamma(S)$, then the fraction of neighbors of v with at least two neighbors in S is at most 2ε. ■

Proof of Claim (1):

Proof: The construction of [24], generates explicitly for any ε, r a right-regular bi-partite graph $G = (L, R; E)$ whose right degree is D_R such that $|L|/|R| = 1 - r$, and for any subset $S \subseteq R$, $|S| \leq |R|\delta$ the neighbor set of S is of size at least $|S|D_R(1-\varepsilon)$, where D_R is the right degree of G. Note that since the right degree is D_R, the average left degree is $D_R|R|/|L| = D_R\frac{1}{1-r}$, which is a constant given that D_R is a constant.

The code is defined by assigning to each left node a parity check over its incident vertices. Let us lower bound the rate of this code: it is at least $r = (|R| - |L|)/|R|$, since each constraint in L removes one dimension of the space. The minimal distance of the code is at least δ, since any non-zero word of weight at most δ is rejected, since there exists at least one check term that "sees" just a single bit at state 1, by Fact (1).

Hence, these are so-called "good" codes. Furthermore, their robustness is at least $1 - 3\varepsilon$ since an error on a set of bits S of size $|S| < \delta n$, is examined by at least $|S|D_R(1-\varepsilon)$ constraints. By Fact (1) at least $1 - 2\varepsilon$ of those constraints, examine S in exactly one location; all constraints that touch a given error set S in exactly one location will be violated; hence the total number of constraints that will be violated is at least $|S|D_R(1-\varepsilon)(1-2\varepsilon) \geq |S|D_R(1-3\varepsilon)$. ■